NAVAL POSTGRADUATE SCHOOL

Monterey, California

THESIS

IMPLEMENTATION CONSIDERATIONS FOR A VIRTUAL PRIVATE NETWORK (VPN) TO ENABLE BROADBAND SECURE REMOTE ACCESS TO THE NAVAL POSTGRADUATE SCHOOL INTRANET

by

Richard Scott Cote'

December 2000

Thesis Co-Advisors: Rex Buddenberg
Daniel Warren

Approved for public release; distribution is unlimited.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.

1. AGENCY USE ONLY ***(Leave blank)***	**2. REPORT DATE** December 2000	**3. REPORT TYPE AND DATES COVERED** Master's Thesis

4. TITLE AND SUBTITLE Implementation Considerations for a Virtual Private Network (VPN) to Enable Broadband Secure Remote Access to the Naval Postgraduate School Intranet	**5. FUNDING NUMBERS**
6. AUTHOR(S) Cote', Richard Scott	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000	**8. PERFORMING ORGANIZATION REPORT NUMBER**
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)	**10. SPONSORING / MONITORING AGENCY REPORT NUMBER**

11. SUPPLEMENTARY NOTES

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.	**12b. DISTRIBUTION CODE**

13. ABSTRACT ***(maximum 200 words)***

As broadband connections to the home become more prevalent, through Digital Subscriber Lines (DSL) and cable modems, students and faculty will desire to access the NPS Intranet via these new means instead of their 56K modems. The introduction of these new technologies will require NPS to re-evaluate how to allow remote access to their internal resources in a secure way, while still allowing for the use of broadband technologies.

This thesis will examine the alternative methods for implementing VPNs, from simple use of Point to Point Protocols (PPP) to high end specialized internet appliances and gateways. Pros and cons of each will be discussed. A mock-up of the schools network will be created to test each of the discussed methods. Final recommendations will be made for a model that can be used by the NPS to implement a VPN. Also discussed will be how that model may be altered to fit other commands throughout the US Navy who desire similar secure remote access to their internal network resources.

14. SUBJECT TERMS Virtual Private Network (VPN), Remote Access, Public Key Infrastructure (PKI), Broadband Access, Computer Security.			**15. NUMBER OF PAGES** 102
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	**18. SECURITY CLASSIFICATION OF THIS PAGE** Unclassified	**19. SECURITY CLASSIFI- CATION OF ABSTRACT** Unclassified	**20. LIMITATION OF ABSTRACT** UL

IMPLEMENTATION CONSIDERATIONS FOR A VIRTUAL PRIVATE NETWORK (VPN) TO ENABLE BROADBAND SECURE REMOTE ACCESS TO THE NAVAL POSTGRADUATE SCHOOL INTRANET

Richard Scott Cote'
Lieutenant, Supply Corps, United States Navy
B.S., State University of New York College at Geneseo, 1990

Submitted in partial fulfillment of the
requirements for the degrees of

MASTER OF SCIENCE IN INFORMATION TECHNOLOGY MANAGEMENT

from the

NAVAL POSTGRADUATE SCHOOL
December 2000

Author: Richard Scott Cote
Richard Scott Cote'

Approved by:
Rex Buddenberg, Thesis Co-Advisor

Dan Warren
Daniel Warren, Thesis Co-Advisor

Dan C Boger
Chairman, Information Systems Academic Group
Dan Boger

ABSTRACT

As broadband connections to the home become more prevalent, through Digital Subscriber Lines (DSL) and cable modems, students and faculty will desire to access the NPS intranet via these new means instead of their 56K modems. The introduction of these new technologies will require NPS to re-evaluate how to allow remote access to their internal resources in a secure way, while still allowing for the use of broadband technologies.

This thesis will examine the alternative methods for implementing VPNs, from simple use of Point to Point Protocols (PPP) to high end specialized internet appliances and gateways. Pros and cons of each will be discussed. A mock-up of the school's network will be created to test each of the discussed methods. Final recommendations will be made for a model that can be used by the NPS to implement a VPN. Also discussed will be how that model may be altered to fit other commands throughout the US Navy who desire similar secure remote access to their internal network resources.

It should be noted that the thesis will concentrate on remote secure access to an internal network from a single remote host more than on the VPNs' additional ability to remotely connect two or more secure networks together, such as can be found in a business to business (B-to-B) environment.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACRONYMS

3DES	Triple-Data Encryption Standard
AH	Authentication Header
BDC	Back-up Domain Controller
CA	Certificate Authority
CBC	Cipher Block Chaining
CHAP	Challenge Handshake Authentication Protocol
CRL	Certificate Revocation List
DES	Data Encryption Standard
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Server or Services
DoD	United States Department of Defense
DON	Department of the Navy
DSL	Digital Subscribers Line
EAP	Extensible Authentication Protocol
ECB	Electronic Code-book Mode
ESP	Encapsulated Security Protocol
FY	Fiscal Year
GRE	Genetic Routing Encapsulation
HMAC	Hashed Message Authentication Code
ICV	Integrity Check Value
IDEA	International Data Encryption Algorithm
IETF	Internet Engineering Task Force
IKE	Internet Key Exchange
IP	Internet Protocol
IPCP	IP Control Protocol
IPSec	Internet Protocol Security
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISAKMP	Internet Security Association and Key Management Protocol
ISP	Internet Service Provider
IT	Information Technology
KDC	Key Distribution Center
L2F	Layer 2 Forwarding protocol
L2TP	Layer 2 Tunneling Protocol
LAC	L2TP Access Concentrator
LCP	Link Control Protocol
LDAP	Lightweight Directory Access Protocol
LNS	L2TP Network Server
MAC	Message Authentication Code or Media Access Control
MD4	Message Digest 4

MD5	Message Digest 5
MIT	Massachusetts Institute of Technology
MPPE	Microsoft's Point-to-Point Encryption
MS-CHAP	Microsoft Challenge Handshake Authentication Protocol
NAS	Network Access Server
NCP	Network Control Protocol
NIST	National Institute of Standards and Technology
NMCI	Navy and Marine Corps Intranet
NOC	Network Operations Center
NPS	Naval Postgraduate School
NSA	National Security Agency
OSI	Open Systems Interconnection
PAP	Password Authentication Protocol
PDC	Primary Domain Controller
PFS	Perfect Forward Security
PGP	Pretty Good Privacy
PKI	Public Key Infrastructure
PPP	Point-to-Point Protocol
PPTP	Point-to-Point Tunneling Protocol
RADIUS	Remote Authentication Dial In User Service
RAS	Remote Access Server or Services
RSA	Rivest, Shamir, and Adleman
SA	Security Association
SAD	Security Association Database
SHA-1	Secure Hash Algorithm 1
SOHO	Small Office Home Office
SPAWAR	Space and Naval Warfare Systems Command
SPD	Security Policy Database
SPI	Security Parameter Index
SSH	Secure Shell
TCP	Transmission Control Protocol
VPN	Virtual Private Network
WAN	Wide Area Network
WWW	World Wide Web
XAUTH	Extended Authentication

ACKNOWLEDGEMENT

The author would like to acknowledge the generous contributions SPAWAR made with their funding, equipment, and time. If not for SPAWAR, and specifically John Feist, the author could not have even dreamed of talking to the leaders in the field of VPN technology, or interacting as he did with cutting edge technology.

The author also wants to recognize certain people, for without whom, this thesis would not be what it is today. First and foremost is his loving wife, the sunshine of his life, for her unwavering support, patience, and understanding. She has always been there when I needed her, even when I couldn't be there for her, and she gladly accepted that, for as she loves to remind me, "The toughest job in the Navy is that of a navy wife."

Thanks also goes out to Rex Buddenberg and Daniel Warren for advising me on my thesis, and for asking the tough questions; to the Network Operations team here at NPS, and specifically to Lonna Sherwin and Bob Gentry for allowing me the use of the new research lab, and Carol Rojas for the use of the equipment; the guys at the Marine Corps NOC in Quantico, including that hard charger GySgt Davis; to Ernie Hernandez, my section leader and close friend for all his help and technical assistance, as well as for the laughs, and Jay Matos for his support; to William Bullier for showing me nothing is impossible, as he builds his second computer at the age of 82; and lastly to Paula Morgan for her support in managing my endless travel plans and claims.

And to Albert Einstein, for his insight, which kept me going when time got tough.

"If we knew what it was we were doing, it would not be called research, would it?"

Albert Einstein

I. INTRODUCTION

A. BACKGROUND

Today's computers have become a replacement for older storage mediums such as notebooks and hardbound encyclopedias. People now store their information on digital media in their desktop computers or on file servers located on an internal network. The convenience of digitally storing this information in a central location has the drawback of not always being immediately available if it is stored on a remote system. A potential solution to this is allowing remote access to this information from another host system, such as a laptop when on travel, or from a person's home computer.

This is a very common situation. Students and faculty at the Naval Postgraduate School (NPS) require such remote access to their digitally stored information, as well as the network internal to the school, herein referred to as the campus intranet. Many professors now post their course wares on the intranet for student access, and ask students to place assignments in designated public folders on the intranet. Many military web sites require the user to be on another military domain (i.e., .mil) to gain complete access to their sites. Being on the campus allows individuals to easily access these resources, but there are times when people are away from the campus and still desire, or even require, access as if they were on the intranet. This could be for something as simple as working on an assignment from home, or as complicated as being on travel in a foreign country and requiring access to important research material. Allowing remote access enables these individuals to seamlessly become part of the intranet.

NPS has met the needs of faculty and students by allowing dial-up access to the intranet. Modem pools allow an individual to dial directly into the network, authenticate themselves, and become part of the intranet. This solution has been well received in the past, but new broadband technologies such as Digital Subscriber Lines (DSL), cable modems, and broadband wireless connections are changing the current equation. These broadband technologies use a new type of modem that cannot be used to dial-in to the intranet via the existing modem pool. They work over the Internet Protocol (IP) and require access to the intranet directly from the Internet. This is not to say that most of these users do not have access to modems and telephone lines. Most users still do, and they still dial into the intranet, but this does not allow them to utilize the larger bandwidth, and therefore faster speeds, afforded them by broadband technologies.

Older modem pools are also limited in a variety of ways. Users are limited by the number of available modems, a problem exemplified by America On Line when they could not meet their expanded user base due to their limited number of modems. Download and upload speeds for information transfer is very limited by both the lower speeds of older analog modems and telephone line conditions. There is also an inherent lack of security of dial in connections, since all data is passed "in the clear" (unencrypted) and can be easily tapped. Finally, there is the issue of *maintaining* a large numbers of modems on the campus, as well as the costs incurred by remote users who must call in via long distance phone calls.

An emerging solution that can be used to address these concerns is a Virtual Private Network (VPN). Used to create encrypted "tunnels" between two hosts, a VPN allows users to access intranets from across the Internet, using their access method of choice, whether it be broadband or dial-up through a local Internet Service Provider (ISP). Connection to the intranet is via the campus' connection to the Internet, eliminating the requirement to connect using the campus analog modem banks. By using *any* access to the Internet, and a VPN, many of these previously discussed limitations are overcome. Data is protected while passing over the Internet by encrypting all the data before it leaves the remote host, and is then decrypted by the local host on the intranet, and vice-versa. Once authenticated, a user can also access the campus intranet just as if they were on campus. They can even be set up with a local NPS IP address, allowing them access to those restricted military domains referenced earlier.

Installation of a VPN on the NPS would also have the added benefit of positioning NPS to become part of the Navy and Marine Corps Intranet (NMCI), a secure Department of the Navy (DON) wide computer network. The NMCI is designed to allow its 450,000 plus members to securely exchange data, voice and video from their desktops (Peniston). This extranet will be based upon a VPN solution, and interconnect both its physical infrastructure and personnel through this method of secure tunneling. With a VPN solution in place, NPS can more easily manage its transition into the NMCI.

B. PURPOSE OF RESEARCH

The intention of this thesis is to create a model for implementing a VPN on the NPS campus. By reviewing emerging VPN technologies, it will discuss methods that allow secure access to the Naval Postgraduate School's Intranet. It will propose a plan encompassing hardware and software solutions, as well as other relevant issues.

Specific research questions the author sets out to answer are:

1. Why do students and faculty desire and/or require remote access to the campus intranet?
2. What are the methods that can be used to gain secure access to the intranet? What are the advantages and disadvantages of each?
3. What are Virtual Private Networks (VPNs) and how can they be used to aid in obtaining secure remote access?
4. Specifically, what are the emerging technologies associated with VPNs, and how are they being implemented?
5. What are the different authentication methods that can be used to authenticate a user? What are the advantages and disadvantages of each?
6. How do IPSec, PKI, and specifically DoD PKI, impact implementing a VPN at the NPS?
7. Can the model developed for the NPS be modified for use throughout the Navy?

C. SCOPE, METHODOLOGY, LIMITS, ASSUMPTIONS

This thesis is limited in its scope. It is meant to create a set of recommendations on how the NPS could implement a VPN. It includes multiple VPN architectures, leaving it up to the NPS to decide which architecture *best* suits their security and accessibility requirements.

It concentrates on the "road warrior" perspective of a single remote host that requires access to a protected intranet via the Internet. It does not cover, nor is it intended to cover, the concept of securely connecting *remote networks* via the Internet to create an extranet. Though this is a significant feature of a VPN, it is not within the scope of this thesis. The thesis is also intended to be a general model for other Naval commands to implement their own remote access features on their VPNs. Though general recommendations will be included, it is expected that other commands may have slightly different architectures; the concepts and recommendations should be extensive enough for each to interpret the work in accordance with their own networks, and adjust as necessary.

The research methodology used in writing this thesis included a combination of methods. Extensive literary research was first performed using the resources listed in the bibliography. Additional education on the subject was obtained from attending the following conferences: DON Public Key Infrastructure (PKI) and Virtual Private Networking Technology Workshop, April 1999; O'Reilly Open Source Software Convention, July 2000; SANS Parliament Hill 2000, Aug 2000; USENIX, August 2000; VPN Con, September 2000; and SANS Network Security 2000, September 2000. While on travel, the author was also able to visit and interview with the Information Technology (IT) staff at Bentley College, where their VPN was illustrated and extensive discussion on authentication methods lead to significant insight into potential solutions. Hands on experience was gained by working with the NPS Network Operations staff in understanding the current network architecture of the NPS intranet, and from creating a custom research lab that was a mock-up of the intranet. This mock-up was then subjected to the installation of different VPN technologies, including a TimeStep© 7520 gateway appliance.

There are certain assumptions the author made in writing this thesis. Though there will be an extensive review of key concepts, it is expected that the reader still has a basic understanding of how networks work. This includes understanding the Internet Protocol (IP) and functions of basic network components such as routers and firewalls.

D. THESIS ORGANIZATION

The author has segregated this thesis into four chapters. Chapters I and II are meant as an introduction to the concepts that will be covered more in-depth later in the thesis. These chapters should be used as background information for those readers unfamiliar with VPNs and the components that go into creating secure connections. Chapter III contains the heart of the thesis, discussing VPN implementation methodologies and specific concerns with each. It is here that pros and cons associated with remote access methods are covered. Chapter IV encompasses the author's recommendations on actual implementation considerations for the NPS. It also discusses conclusions and possible areas for further research.

II. BACKGROUND INFORMATION AND KEY CONCEPTS REVIEW

A. VPN OVERVIEW

Virtual private networks are used to create secure connections between two hosts over an unsecured medium. This medium can be a public medium such as the Internet, or a private medium, such as a company's Wide Area Network (WAN). Its purpose can be to allow users access to information that normally would not be available to them since they are not on the internal network, or to secure communication on the network they are sharing. When VPNs are used to gain access to a network, they enable you to act as if you are actually on that network, giving you the same access to resources you would normally have if you were there. When used on an internal network, a VPN ensures privacy across the internal network by encrypting data between specific hosts or network segments.

There are three varieties of VPN architecture: network-to-network, host-to-host, and host-to-network (Brenton and Elfering, p. 7). All three are based upon the same

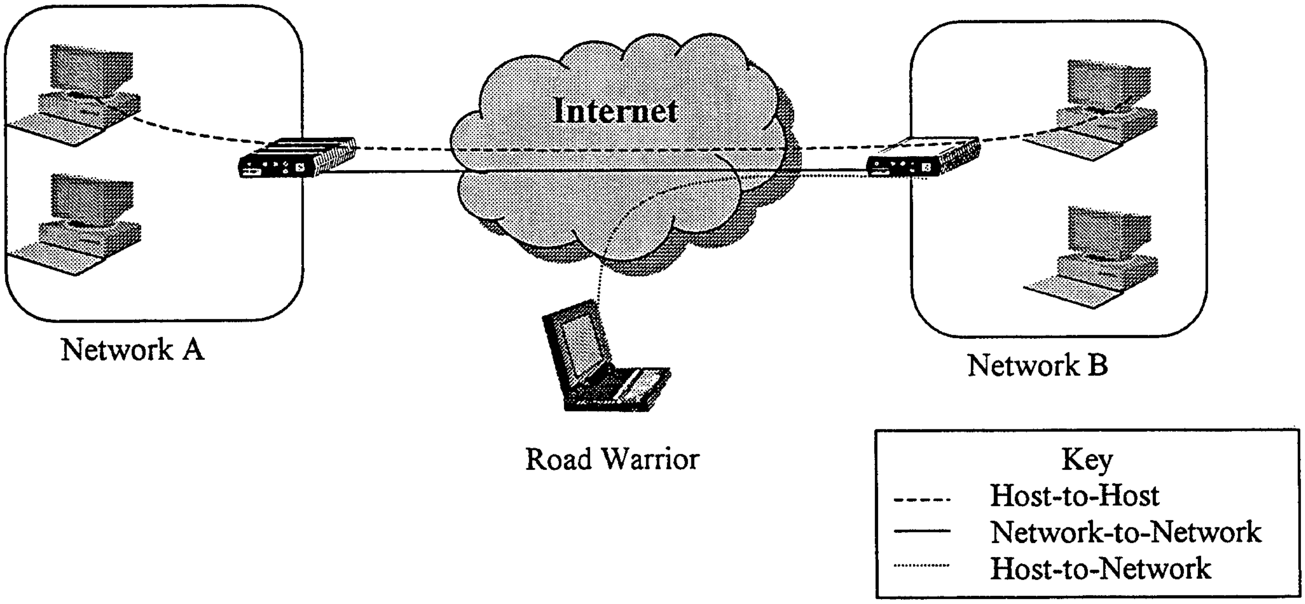

Figure 2-1. Varieties of VPNs

technology, the difference is solely in how that technology is applied.

VPNs can be used to create an *extranet*. This is where a private network is selectively opened to designated parties outside the network. This access to the network is normally privately held among specific members of a firm or institution, where the information is proprietary and closely held (Maier, pp. 6-8). This thesis concentrates on the concept of creating an extranet, specifically in the form of a host-to-network secure VPN. This type of extranet is commonly referred to as the "Road Warrior" scenario in tribute to those members of an organization who perform their work away from the office, but yet still require access to those resources found on the company's intranet.

VPNs should not be confused with WANs and Remote Access Services (RAS). VPNs are similar to WANs and RAS in that they connect remote users to private network, but there are differences that should be pointed out. Both WANs and RAS are more of a physical connection mechanism, where a VPN is more of a logical use of a physical connection. WANs normally consist of two or more networks connected via dedicated and private lines supplied by a third party, such as the telephone company; VPNs utilize an existing connected public medium to create logical tunnels instead, encrypting IP datagrams which may traverse the same connections a WAN may reside on. Where leased lines may give the illusion of security, for one can never be truly sure where your data traverses once it leaves your site, VPNs use their encryption ability to ensure security over any medium. Traditional RAS services are similarly composed of banks of modems using incoming, on-demand telephone lines for connecting remote users, again leaving the illusion of security; VPNs connect via an already established public network access point, such as through a router or firewall, ensuring confidentiality via encryption. Instead of dialing in to a central RAS point, the users of a VPN dial their own Internet Service Provider (ISP) and connect to the network using their secure method or protocol, such as SSH, PPTP, L2TP, etc. (Erwin, Scott and Wolf, pp. 45-46)

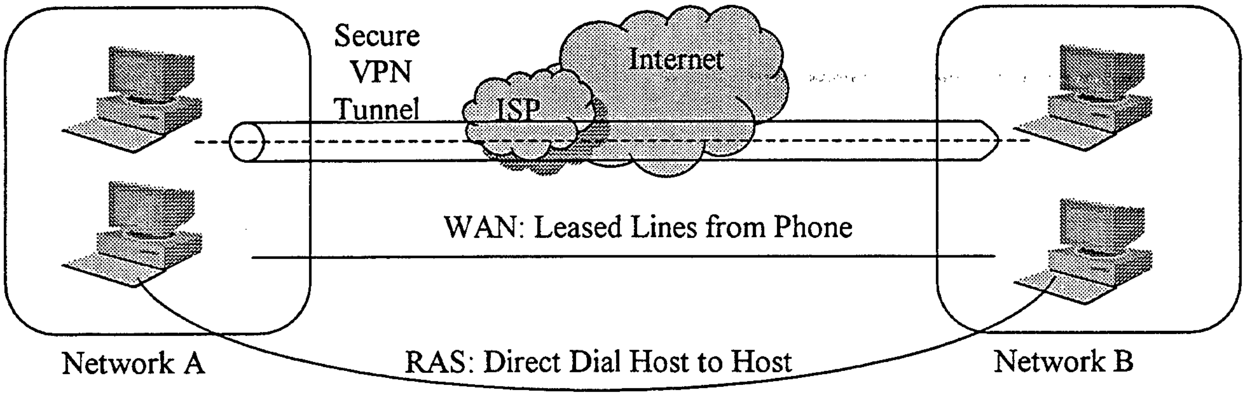

Figure 2-2. VPN vs. WAN vs. RAS

Encryption and authentication methods are combined in a VPN to establish a confidential communication "tunnel." These methods ensure that the properties of authentication, confidentiality, and integrity are met. Each of these aspects is developed in more detail later in the chapter.

The beneficial characteristics of VPNs have lead to early adoption of the technology by many companies and to projections of its substantial growth. Such benefits include the flexibility of being able to quickly create tunnels to connect networks, verse having to wait to install leased lines, as well as not having to pay for their lease. VPNs also benefit from the confidentiality gained through the encryption of its data, versus the *illusion* of confidentiality found with leased lines, as well as the ability to secure data while it transits even an internal network. It is projected that companies can expect savings of 20% to 47% of their WAN costs through releasing leased lines and 60% to 80% of corporate costs for remote access dial-up, recouping monthly charges for 1-800 phone numbers and other phone charges for remote access (Bourne, p. 2). VPNs also allow for the use of private, non-routable IP addresses across the Internet, as well as allowing legacy equipment to operate over extranets with non-routable protocols, through encapsulation. These technologies will allow the average number of telecommuters to grow by 184% between 1999 and 2001, as well as a growth in the number of mobile workers by 72% in the same time frame (Newbridge). Chris Brenton in his VPN and Remote Access book states that VPN technology is expected to expand 300-1000% by 2003, and Newbridge projects that VPNs will grow at a compound annual growth rate of $32 billion in the same period.

There are additional benefits for the military as well. For such organizations as the United States Navy, the VPN's ability to ensure object level security of the data gives the added benefit of no longer having to find ways to keep the network backbone secure. These organizations have been forced to pay an exorbitant amount of money to ensure their private networks remain private, and have had to adopt an assortment of implementation methods in doing so. VPNs can aid in allowing them to consolidate these solutions by investing in 100% insecure, off-the-shelf IP "plumbing" for all communications, including such expensive systems as ship-to-shore communications, and concentrate on securing the bits of data, not the physical, link-level infrastructure.

B. BASICS OF ENCRYPTION, AUTHENTICATION, AND INTEGRITY

1. Confidentiality through Encryption

The concept of the VPN uses encryption to provide communication confidentiality. Encryption is the conversion of data into a protected form for transmission over an untrusted medium to a trusted party. Encryption is meant to be a computationally easy conversion of data to and from cipher-text (encrypted data) by the trusted parties, but computationally difficult for an untrusted party who intercepts the data. Note the subtlety in this concept: encryption is <u>not</u> meant to be impossible to break, but is meant to be too hard or take too long to get at what is hidden in the cipher-text. Encryption only needs to be strong enough to keep the data protected for the lifetime of the *value* of the data. Once the data no longer has value, the encryption of it become unnecessary. (Erwin, Scott, and Wolfe, pp. 22-23)

2. Authentication

Before establishing a secure connection, VPNs strive to authenticate both hosts that are attempting to create the tunnel. This is like asking both ends to "log-in" with each other, usually using a two tier system: what you have and what you know. Each end validates itself by having something unique, such as a piece of client software with an embedded shared secret key or maybe a digital certificate, and also with something they know, such as an additional user ID and password or pass-phrase. More specifics of this concept are given later in this chapter, but it is important to the general concept. By authenticating, both ends ensure that no one untrusted can connect to the network. This keeps intruders out, and helps prevent session hijacking (when a third party takes over your session, like a terrorist hijacking a plane) by revalidating users throughout the course of an established session (Brenton and Elfering, p. 46).

3. Integrity

Even if a VPN is utilizing encryption, and both parties have been authenticated, you may still require positive proof that the original data being sent as cipher-text was not changed. Integrity can be designed into your VPN solution through the use of hashing and encryption, all of which is covered in more depth in the Encryption Algorithms section of this chapter. When integrity is enforced, it can also aid in the process of non-repudiation. This is when authentication, integrity, and encryption are combined into a method that proves a message was sent by a specific person, and was not altered in any way.

C. VPNS AS A SYSTEM

In its most basic form, a VPN is the use of mathematical algorithms by a set of procedures in an end product's implementation. Cryptographic algorithms are used to convert data via encryption and in a process called hashing. These algorithms are then used in a set of procedures, such as PAP and IKE, to create authentication and integrity of the data. These procedures are then standardized for their implementation, as found in the protocols PPP, PPTP, L2TP, and IPSec. Finally, these implementations are placed into a company's end products, the VPNs themselves, and take a variety of forms such as Microsoft's PPTP implementation, embedded VPN products within vendors' firewall products, and gateway appliances such as Alcatel's TimeStep product line. Each of these areas is covered in much more depth throughout the rest of this thesis.

1. Tunneling

Tunneling is the common term for encapsulating individual packets on a packet-switched network and is the basis upon which VPN implementations are based. It takes the original packet and encapsulates it into a new packet with new header information, making the original packet the payload of the new packet. Tunneling is extremely useful for carrying non-routable protocols such as NetBIOS over a WAN using TCP/IP. In the case of IPSec, IP packets are being tunneled in new IP packets for protection, if encrypted, or possibly to use non-routable IP addresses over a public network. There is tremendous potential in using private IP address spaces, such as 192.168.X.X and 10.X.X.X, since all that is required is one public address used at the gateway. With IP addresses becoming scarce in IP Version 4 (IPv4), and most home users with broadband access using private addresses to share their Internet connections in their homes, this type of encapsulation will become quite common. A caveat to this is that Internet Protocol Version 6 (IPv6), if ever adopted, has been crafted to allow for IPv4 encapsulation and vice versa, as well as an expanded IP address space of "one undecillion addresses [translating] to more than a thousand IPv6 addresses for every square meter on the surface of the Earth" (Walton, p. 6). The IPSec hosts usually perform the encapsulation, whether it is the laptop of the road-warrior, or the gateway of a network. Additional information on IPSec tunneling can be found in RFC 2003 *IP Encapsulation within IP*. (Murhammer, et. al., pp. 49-50)

The following sections go into greater detail on cryptographic algorithms, the frameworks they are used within, and the implementations of IPSec and authentication. They are meant to develop a greater understanding of the VPN as a system, giving the reader the building blocks for understanding VPN protocols and methodologies covered in later chapters of this thesis.

D. ENCRYPTION ALGORITHMS

In its simplest form, a cryptographic algorithm is a "procedure that takes the plain text data and transforms it into cipher-text in a reversible way" (Snyder, p. 13).

There are two methods used to encrypt data. The first is based upon a secret shared or common key. This method is called symmetric key cryptography. The second is based upon a mathematical algorithm that allows for the generation of two keys; each only able to decrypt messages encrypted by its mate. This latter method is the basis for many procedures found in VPN implementations in use today, and is called asymmetric key cryptography.

Characteristics of a good cryptographic algorithm include:

- The algorithm is known and subject to analysis and no practical weaknesses have been discovered
- Given the clear-text (input to the algorithm) and the cipher-text (output of the algorithm), the key can still not be computed

If the algorithm is published or known, not secret, it provides stronger encryption than "security through obscurity" where it is assumed just because an algorithm is not known, it is secure (Snyder, pp. 20-21). The encryption itself depends on two mechanisms: the encryption algorithm, which provides the mathematical conversion mechanism, and a key, which provides the randomness to the final cipher-text when used by the encryption algorithm (Bird, p. 22).

1. Symmetric Key Cryptography

Symmetric keys have long been used throughout history, starting from the time of Caesar. This method of a shared secret key has even been at the foundation of our military security since the military's introduction of cryptography. As its name implies, the system uses a single key shared between two hosts, and the users of the system closely hold the *only* key that can decrypt and encrypt. To be used, both parties must agree on the key *before* any secure communication can take place between (Murhammer, et. al., p. 29). Therefore, a system must be developed to distribute and protect these shared secret keys.

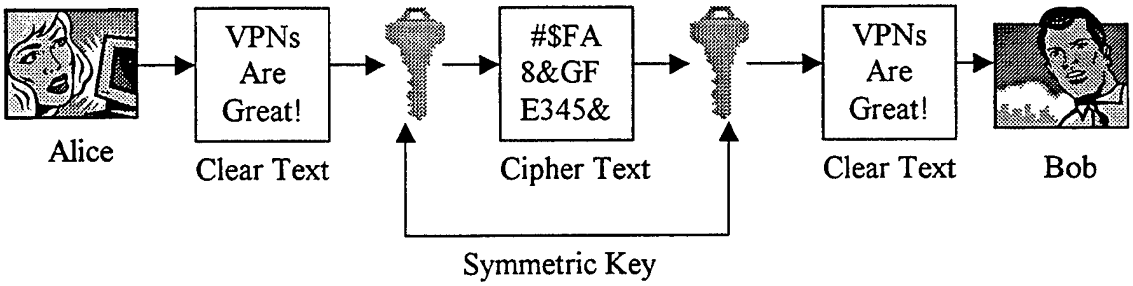

Figure 2-3. Symmetric Key Encryption/Decryption

The most common types of symmetric key algorithms are block algorithms, which work on blocks of bits of the clear-text message at a time. Block ciphers modes include Electronic Code-book Mode (ECB), where each block of cipher-text is encrypted independently, and Cipher Block Chaining (CBC), where the result of the previous block is used in the encryption of the next block. Well known block ciphers include: Data Encryption Standard (DES) with a 56 bit key length; triple-DES (3DES), where DES is applied three times to the clear-text; and International Data Encryption Algorithm (IDEA) which uses 64 bit blocks and 128 bit keys. IDEA is most notably used in the software application Pretty Good Privacy (PGP). Though DES is an older standard, developed in the 1970's, it is still widely used today. (Murhammer, et. al., pp. 29 - 30)

2. Asymmetric Key Cryptography

Where symmetric key cryptography requires sharing a single key that must be kept secret from non-trusted parties, asymmetric key cryptography allows users to generate their own pair of mathematically related, yet individual, keys. Then one of these keys (called the public key) is shared publicly and used to encrypt messages back to them by any other party. They keep their other key (called the private key) private only to themselves and use it to decrypt any messages sent to them. Since the keys are complementary, but not identical, they cannot be used to perform bi-directional encryption/decryption, meaning once something is encrypted with one key it can not be decrypted by that same key; only its mate can decrypt it. This is important because once a message is encrypted with a publicly shared key, no one, not even some other party that has a copy of the public key, can decrypt the message. An encrypted message is then safe to transmit over insecure or untrusted mediums; even if an untrusted party intercepts the message, that party cannot decrypt the cipher-text with the publicly available key.

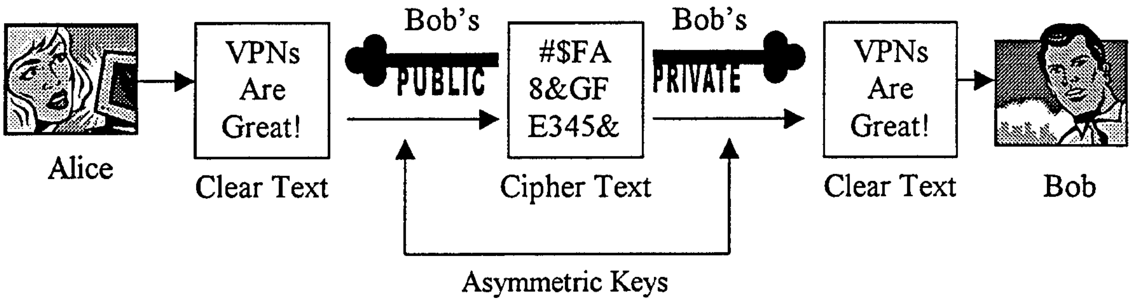

Figure 2-4. Asymmetric Key Encryption/Decryption

3. Hashes

A hash is a one-way algorithm used to create a unique fixed length cipher-text of a message, or more simply put, it's a formula to convert a message of any length into a single string of characters called a message digest. It's considered a zero-key encryption algorithm; where symmetric key encryption uses a single key for encryption/decryption, and asymmetric key encryption uses a second key for decryption, a hash cannot be decrypted, so it has no key. Think of it in terms of mashing a potato: once a potato is mashed, reconstructing the original potato is rather difficult.

Hashes are mainly used as a fingerprint of the clear-text. Hashing algorithms are made to be collision-resistant, meaning it's highly improbable that two different clear-text messages will produce the same hashed value. This is what provides for message integrity as discussed earlier. When a hash of the clear-text is sent with the encrypted message, the message once decrypted can be re-hashed with the same algorithm, and the hashes compared. If they are the same, then the message has not changed.

An extension to the concept of hashing is Message Authentication Code (MAC). This is when you either concatenate a key to the clear-text before hashing it, or when the hash algorithm can take a key as a second input (the message being the first input) to the hash. Simply put, if you encrypt a hash, it becomes a MAC; if you concatenate a secret key to a message and then hash it, it becomes a MAC (Murhammer, et. al., p. 37). A sender utilizes a MAC hash by including the session key with the message when hashed, so when the receiver rehashes the concatenation of the message with the session key, it ensures authentication when the two match. The greatest value of a MAC is its ability to provide message integrity, without the intensive computation required to encrypt a whole message.

Hashes can also play a role in digital signatures. This is where the hash of the original message is encrypted with the private key of the sender. Once the message is received, the public key is used to decrypt and get the original hash value and is then compared to a new hash of the decrypted message. Since the encrypted hash can only be decrypted with the public key of the original sender, and only the sender has access to their own private key, it ensures that sender is who they say they are.

The most widely used hash functions are MD5, created by Ron Rivest, and Secure Hash Algorithm 1 (SHA-1), adopted by the National Institute of Standards and Technology (NIST) and the National Security Agency (NSA). MD5 produces a 128-bit fixed length hash and SHA-1 produces a 160-bit fixed length hash. It should be noted that neither of these takes a key as an input parameter, and therefore cannot be used for MAC calculations, unless the key is concatenated to the message.

Hashed Message Authentication Code (HMAC) makes a hash function a MAC. It applies the hash twice in succession. Here, MD5 can be used on a concatenated message/key combination, then the results of that hash are concatenated again with the key, and hashed a second time. (Murhammer, et. al., pp. 29 - 30)

E. PROCEDURES AND FRAMEWORKS FOR USING ENCRYPTION ALGORITHMS

1. Diffie-Hellman

Diffie-Hellman utilizes asymmetric keys (i.e., one public, one private), to create a unique symmetric key. It is used extensively in key creation in VPNs, and is worth further explanation here.

Diffie-Hellman allows two hosts who share no common keys to create a shared secret key by using asymmetric keys. It sounds complicated, but really isn't. The following diagram and detailing of steps will clarify it.

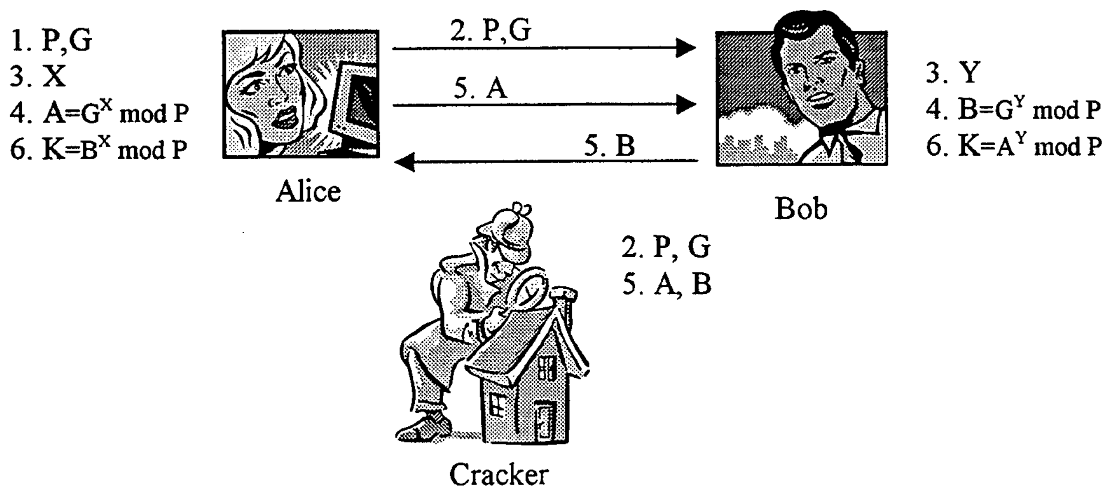

Figure 2-5. Diffie-Hellman Key Creation

This diagram shows how once two hosts have established communications, they can, with no prior shared information, exchange data that is susceptible to interception and still create a shared secret that cannot be reversed engineered by the information that was intercepted. It starts in step one by host one, Alice, choosing a large prime number P and an integer G. In step two, those numbers are shared with host two, Bob, and since they traverse the public medium, we assume bad guy Cracker can see them as well. The third step is for both Alice and Bob to chose a random number. Alice chooses X, and Bob chooses Y. These numbers are their "Private keys" and not shared. In the fourth step, both Alice and Bob compute their "Public keys" by raising integer G to the power of their private key and performing modulus arithmetic on the product with the large prime number P. So Alice's public key is $A=G^X$ mod P, and Bob's public key is $B=G^Y$ mod P. In step five, Alice sends Bob her public key A, and Bob sends his public key B to Alice, and since it again traverses a public network, we assume Cracker gets these as well. At this point Cracker has P, G, A, and B. The final step to creating the symmetric shared secret key is to compute it by raising the received public key by the power of the private key and taking the modulus with P. So Alice computes the shared secret key K by computing $K=B^X$ mod P, and Bob computes the same shared secret key K by computing $K=A^Y$ mod P. Now they have a shared key that they can encrypt and decrypt messages with (since it is a symmetric key), and Cracker can not crack it, since it is computationally infeasible to compute both Alice's and Bob's private keys, even given the information Cracker intercepted.

The reason this works is due to the simple mathematical principles that $((G)^X)^Y=((G)^Y)^X$ and $K=(G)^{XY}$ mod P (remember, X and Y are the private keys that never traverse the untrusted medium).

2. Public Key Infrastructure (PKI)

PKI is the infrastructure used to create, distribute, and revoke the public key of an asymmetric key pair by a central authority, called a Certificate Authority (CA). The CA is responsible for creating public key certificates for each person who uses the CA's service. Each certificate includes not only a user's public key, but also some form of distinguished name to uniquely identify the individual, a validity period, and some form of the CA's digital signature that is used to verify that the certificate is authentic (Snyder, p. 39). Though a certificate may sound imposing, it can be thought of being similar to a driver's license or passport; it's a form issued to you by a trusted authority that can aid in validating who you are, similar to your license and passport issued by your state or federal government.

It should be noted that for smaller VPN implementations, the full-blown infrastructure of a PKI in not necessarily required. Pretty Good Privacy (PGP) is an example of an asymmetric key infrastructure that doesn't require a CA, but can still be used for encryption and authentication. The use of PGP in this way assumes a secure method is established for issuing public keys, such as handing a copy of a public key on disk to a co-worker, so they know it is from you. In this case, the manageability is very similar to using local host files instead of a Domain Name Service (DNS) for correlating a computer Media Access Control (MAC) address to an IP address; its fine in a very small implementation, but the overhead becomes non-trivial when needed to scale to larger implementations.

3. Internet Key Exchange (IKE)

In order for hosts to create the secure tunnels that make up a VPN, they must create secure associations, authenticate each other, and exchange the keys they will use for the encryption of their data. This is done with the Internet Key Exchange (IKE) protocol. IKE is responsible for providing authentication of all peers, handling the security policy negotiations, and controlling the exchange of keys (Scott, Wolfe and Erwin, p. 36). It is used to create an initial encrypted session, enabling the exchange of the information required to make the final encrypted session.

This protocol originates from the combination of two other protocols: the Internet Security Association and Key Management Protocol (ISAKMP) and Oakley. IKE inherits its security association (SA) and key management (but not key exchange) from the ISAKMP protocol, and supports the pre-shared key, digital signature, and public key encryption methods of authentication. Oakley was originally the key exchange protocol used for VPNs and lent its abilities to IKE. Oakley is vital to security, being that no matter how strong the encryption and authentication algorithms are, they are worthless if your key is compromised. (Murhammer, et. al., p. 71)

IKE is performed in two phases: phase one and phase two. In phase one, two hosts establish a secure connection, called the IKE SA. Authentication is either incorporated into phase one with digital certificates, or takes place between phases one and two, kind of phase 1.5, with extended authentication techniques (Bird, pp. 89-90). In phase two, the final keys used for encryption will be generated, and the IPSec SA will be negotiated. After these two phases, the IPSec SA has been created and is used for all further communication, thereby making the secure "tunnel."

In the first phase, IKE chooses between two modes to complete the IKE SA: Main mode and Aggressive mode. During the second phase, IKE uses its third mode, Quick mode, to negotiate the final IPSec SA. If extended authentication is required, another step is inserted between these two to authenticate the user.

a. Main Mode

To create an IKE SA, Main mode will take the two hosts through three two-way exchanges, totaling six steps. The first exchange, being steps one and two, the hosts agree on basic algorithms and hashes that will be used throughout the remaining four steps. In steps three and four, they prepare to create an encrypted tunnel by using Diffie-Hellman, exchanging the necessary items to create their shared secret key, as well as their digital certificates if they have them. Steps five and six complete the Diffie-Hellman exchange, leaving each with the shared secret with which to encrypt the rest of the data that they will share to finish the IKE and create the IPSec SA. (TimeStep, pp. 20-21)

b. Aggressive Mode

This mode is used to accomplish the same end result as Main mode, but it does so in only three steps instead of six. This mode is quicker, but at the sacrifice of not protecting the identity of each host that is not normally divulged until the encrypted fifth and sixth steps of Main mode. This is accomplished by sending the IKE SA proposal, the Diffie-Hellman information, digital certificates if used, and the ID packet all in the initial exchanges between the two hosts during steps one and two, instead of breaking them down into the six steps. The last step of this mode is just a confirmation exchange. (TimeStep, p. 21)

c. *Extended Authentication (XAUTH)*

If the hosts chose not to use digital certificates to authenticate each other, then authentication must take place before continuing on to phase two, Quick mode. Here, an extra set of exchanges takes place between the hosts to authenticate each other using an extended authentication method such as RADIUS. For more information on authentication, see the Implementation of Authentication section later in this chapter.

d. *Quick Mode*

Quick mode is used in phase two to negotiate the final IPSec information between the now-authenticated and secure hosts. This is accomplished using the current IKE SA to ensure security of the exchange. The end result of this phase is the creation of the final IPSec SA and the generation of fresh keying material, including the symmetric key used for all further encryption between the hosts.

To generate the new SA, the initiating host sends the Quick mode message, protected by the IKE SA, requesting the new IPSec SA. This request includes which Security Parameter Index (SPI) to use in future communications to it. This SPI, combined with the destination IP address and protocol to be used, uniquely identifies a single IPSec SA. It is important to remember that these SAs are only for one-side of the conversation; it takes two SAs for bi-directional communication. Security Parameter

Indexes are covered in greater detail in the next section of this thesis, Implementation of the IPSec protocol.

e. *Perfect Forward Security (PFS)*

PFS is an option in generating the final set of symmetric keys in Quick mode. It is used to create a key that does not have any information derived from or depending on the previous key used.

When the final keys are to be created, there are two ways this can be done. Quick mode can be set to create the new key by just hashing the original key used during IKE phase one. Though this is simple and fast, the problem is that if your original key is later cracked, it is a simple step to hash that cracked key, giving the cracker your next key. To keep this from happening, Quick mode can use PFS to initiate a new Diffie-Hellman key exchange to create a new key each time one is needed. This use of Diffie-Hellman ensures that even if the original key is cracked, it gives no information on the next key. The cracker would then have to go back and crack the new key instead of just being able to hash the previous key value to find the new key. The down side to this is that it takes more time, steps, and computational power to create new Diffie-Hellman keys then is does to just hash the previous one, so security needs to be balanced with performance to meet each individual communities needs. (TimeStep, pp. 19-23)

F. IMPLEMENTATIONS OF THE INTERNET SECURITY PROTOCOL (IPSEC)

The IPSec protocol is a suite of protocols combined to create a security standard by the Internet Engineering Task Force (IETF). Its purpose is to act as an extension to the Internet Protocol (IP), providing security services, since IP was not designed with security in mind. One of the goals of the IETF was to create an implementation that would be an interoperable, vender neutral standard. It also provides a framework that decreases implementation flaws by allowing for the adoption of new implementations that are compliant with the standards. The use of a framework like this allows independence from specific algorithms, locking system thinking into the standard.

It consists of three components, including encryption algorithms, key management procedures, and authentication implementations. These general components define the architecture of the protocol, making it so that new algorithms and methods can be added to the framework with very little work and having little effect upon previous implementations. (Erwin, Scott, and Wolfe, p. 33)

IPSec supports two modes, transport mode and tunnel mode; consists of two protocols, the ESP protocol and the AH protocol; and uses the Internet Key Exchange (IKE) to determine authentication methods, security associations, and to agree upon key generation and key rotation techniques. (Erwin, Scott, and Wolfe, p. 33)

1. Transport Mode

Transport mode consists of securing a packet's payload through one of the two security protocols, while leaving its packet's IP header information untouched. An IPSec header is placed inside of the IP datagram packet, after the original IP header but before the payload, and the payload can be either encrypted or in the clear, depending on the security protocol chosen.

An example of when this type of mode would be used is when information on an internal network requires authentication or encryption, but capturing the source and destination addresses do not give away any vital information, since it is internal to the network anyway. If you wanted to protect those machine addresses, say across a public untrusted network (so an interceptor can not determine the address of your mail server), you would use tunnel mode.

2. Tunnel Mode

Tunnel mode encapsulates the original packet as data within another packet. The source and destination addresses in the outer packet correspond to the VPN device involved, not the specific host. This encapsulation secures not only the packet payload, such as done in transport mode, but encapsulates the entire packet with a new header. Again, the packet can be either encrypted or in the clear, depending on the security protocol chosen.

An example of when tunnel mode would be a preferred choice is when you are using private, non-routable addresses over a WAN, or you want to protect your internal address information when a packet transits a public, untrusted network.

3. IP Protocol 50 Encapsulating Security Payload (ESP)

ESP can be used to provide encryption, integrity checking, and/or authentication to an IP datagram. The set of desired services is selected upon negotiation of the SA. It is performed on each individual packet, on a per-packet basis. Both the integrity and encryption aspects are optional, and either or both can be selected for use. If both are selected, the receiver of the packet must first authenticate the packet, and only if it authenticates, will the packet be decrypted. Integrity check and authentication methods must be chosen so that they complement each other. The encryption algorithm used is independent of whatever choice is made for integrity and authentication. (Murhammer, et. al., pp. 56-57)

4. IP Protocol 51 Authentication Header (AH)

AH is used to provide integrity and authentication, but not encryption, to an IP datagram. It is used to authenticate the packet source and ensure that the packet has not been altered during transmission. AH is performed on each individual packet, on a per-packet basis. The data authentication resides in the Integrity Check Value (ICV), inside the AH Header. It is produced using the algorithm negotiated during the SA, and is usually one of the following: HMAC-MD5-96, HMAC-SHA-1-96, Keyed MD5, or Keyed SHA-1. (Murhammer, et. al., pp. 51-53)

Encapsulation Mode Security Protocol	Transport Mode	Tunnel Mode
ESP	Encrypts only the IP packet's payload, places an ESP header between the untouched IP header and the encrypted payload and/or authenticates the packet	Encrypts the entire original IP packet, places an ESP header between the new IP header and the encrypted original IP packet and/or authenticates the packet
AH	Authenticates the entire IP packet, places a AH header between the untouched IP header and the unencrypted payload	Authenticates the entire original IP packet, places an AH header between the new IP header and the unencrypted original IP packet

Table 2-1. Matrix of Encapsulation Modes and Security Protocols

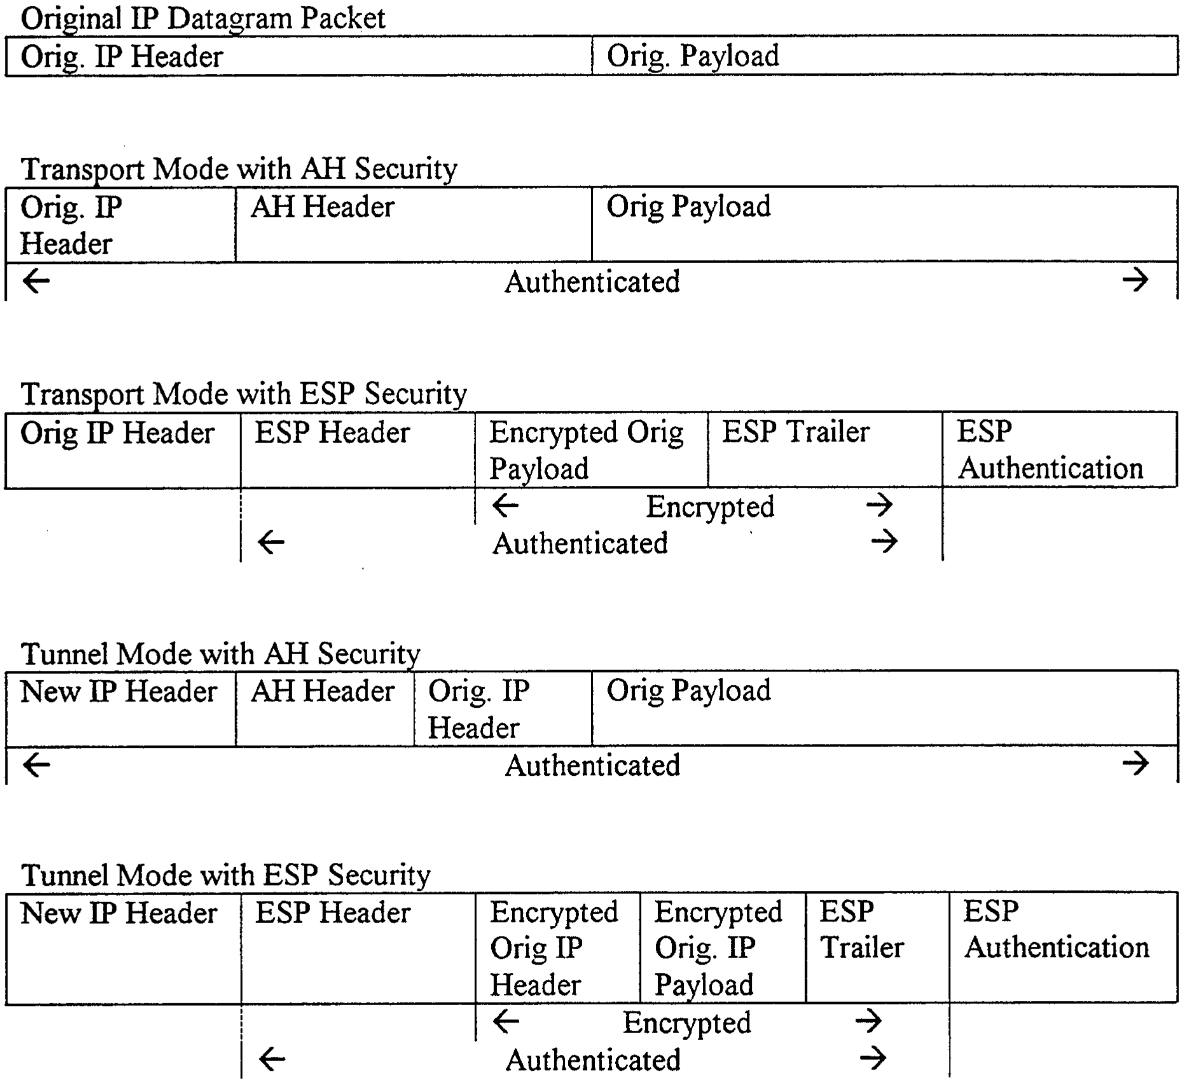

Original IP Datagram Packet

Orig. IP Header	Orig. Payload

Transport Mode with AH Security

Orig. IP Header	AH Header	Orig Payload
←	Authenticated	→

Transport Mode with ESP Security

Orig IP Header	ESP Header	Encrypted Orig Payload	ESP Trailer	ESP Authentication
		← Encrypted	→	
	←	Authenticated	→	

Tunnel Mode with AH Security

New IP Header	AH Header	Orig. IP Header	Orig Payload
←		Authenticated	→

Tunnel Mode with ESP Security

New IP Header	ESP Header	Encrypted Orig IP Header	Encrypted Orig. IP Payload	ESP Trailer	ESP Authentication
		←	Encrypted	→	
	←	Authenticated		→	

Figure 2-6. Effects of Encapsulation Modes and Security Protocols on IP Packets (Murhammer, et. al., and pp. 54-61).

5. Security Associations (SA)

An SA is a logical security connection between two IPSec systems. SAs are negotiated when IPSec hosts create their connection, and take the form of:

<Security Parameter Index, IP Destination Address, Security Protocol>

a. Security Parameter Index (SPI)

A unique 32-bit value used to identify the SA. It is carried in the header of the security protocol, and selected by the destination system during SA establishment. (Murhammer, et. al., p. 48)

b. IP Destination Address

Is the IP address of the destination host.

c. Security Protocol

Either ESP or AH.

SAs define a security association in only one direction, to the destination. Since communication is a bi-directional activity, two SAs must be established during an IPSec session, one in each direction. (Murhammer, et. al., p. 48)

6. SA Bundle

Since an SA can only be based upon a single Security Protocol, ESP or AH, if the IPSec session requires both, an SA Bundle must be created. SA Bundles are two SAs defined in each direction, establishing a total of four SAs for the IPSec session. This is seen most often when a mobile host needs to create an AH SA between itself and the network gateway, and a nested ESP SA extends to the host behind the gateway. (Murhammer, et. al., p. 48)

7. Security Policy Database (SPD)

SPD is a database of the available security policies that can be used to negotiate an SA. It contains an ordered list of policy entries that two hosts will compare to determine what security policy they will use to establish a connection. The SPD must specify the security services provided, protocols employed, algorithms used, etc. It can be thought of being similar to a rule base for a packet filtering firewall or an access control list in a router. (Murhammer, et. al., p. 48)

8. Security Association Database (SAD)

The SAD is a listing of all active SAs. It contains the parameter information about each SA, such as the Security Protocol, SPI, protocol mode, and the SA lifetime. (Murhammer, et. al., p. 49)

G. IMPLEMENTATIONS OF AUTHENTICATION

Part of the IPSec protocol suite is the support for authentication. Ensuring a VPN validates a host is actually a vital part of the process. So the protocol has been designed to support multiple validation procedures. Also note that this discussion is about validating a both hosts and people. This is because a VPN can be set up between two hosts, such as gateway VPN appliances, that have no "person" attached, but yet their identity must also be authenticated.

In the implementation of IPSec, authentication takes place while the final VPN session tunnel is being set up by the IKE exchange. It takes place either with digital certificates during phase one of the IKE exchange, or between phases one and two with extended authentication, XAUTH. It is also at this stage where if authentication fails, the tunnel is broken and the final SAs are not created.

There are quite a few authentication techniques available, and the list is changing all the time, so for the purpose of this thesis, the following four frameworks will be considered main stream and covered: X.509 digital certificates; and the extended authentication procedures LDAP, RADIUS, and Kerberos.

1. ISO X.509 Digital Certificates

This standard is the most highly recommended implementation for authentication in a VPN. It is also the most complicated to manage and implement.

A trusted third party called a Certificate Authority (CA) is responsible for issuing, validating, and revoking certificates for individuals or hosts. These certificates act like passports, aiding, with the use of your signature via your private key, in the validation of who you are. Each certificate is unique, containing a users public key, some form of distinguished name to uniquely identify the individual, a validity period, and some form of the CA's digital signature (using the CA private key) which is used to verify the certificate is authentic (Snyder, p. 39).

The actual authentication takes place when one host, we'll call him Bob, receives something from the other host, who we will call Alice, that has been digitally signed by being encrypted with Alice's private key. Once Bob decrypts the message with Alice's public key, he knows its from Alice, and knows its Alice's public key, because that key has been certified by the CA. Alice is now authenticated to Bob, and the procedure must be done in the opposite direction to authenticate Bob to Alice.

2. Lightweight Directory Access Protocol (LDAP)

LDAP is type of phone book white pages used to locate specific host information, including host identity and authentication information. It is lighter subset of the X.500 Directory Service protocol. Its strength lies in its ability to be easily searched for the host you are looking for, and its ability to catalog and cross-reference hosts by logical identities like business organization or geographic location. It is being widely implemented natively in many key products, including Netscape, Microsoft's Active Directory, and even in Novell's and Cisco's products. (Whatis?com, "Lightweight Directory Access Protocol")

3. Remote Authentication Dial In User Service (RADIUS)

RADIUS is a protocol for exchanging information between a RADIUS server and a RADIUS client. Its main purpose is to provide authentication, authorization, and accounting for remote access users when they connect to a network. It can act as middleware for older authentication techniques, allowing for the continued use of legacy authentication systems.

When used with a VPN, RADIUS either authenticates hosts using its own authentication services, or acts as a middleman between the VPN gateway and another authentication server such as Microsoft's Primary Domain Controller. In the later situation, the VPN gateway informs the RADIUS server it needs to authenticate a user. The RADIUS server then tells the VPN what to ask for (e.g., user name and password), and how to send that information back (e.g., hash the password). The RADIUS server then passes that information onto the domain controller. If the information validates the user (i.e., the user name and hash of the password corresponds to what is in the domain controller's security database), the domain controller informs the RADIUS server that the

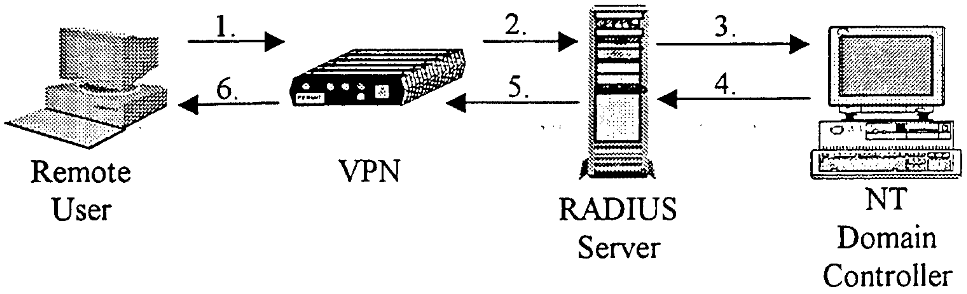

1. Remote user Initiates connection to the VPN

2. VPN request authentication from RADIUS server

3. RADIUS sever requests authentication from Domain Controller

6. IKE continues into Phase Two Quick Mode

5. RADIUS returns authentication to VPN Gateway

4. NT Domain Controller authenticates user

Figure 2-7. XAUTH Authentication with RADIUS Server and NT Domain Controller

client has authenticated. The RADIUS server then informs the VPN that the user has authenticated, and the VPN continues with the IPSec process, going into Quick mode.

4. Kerberos

Kerberos is a network authentication protocol created by the Massachusetts Institute of Technology (MIT). It provides secure authentication based upon a principle's (i.e., user's) knowledge of a password that the user must present. These passwords are stored on the Kerberos server

Kerberos is based upon a shared secret cryptography that is used to authenticate a host to a Kerberos server, called a Key Distribution Center or KDC for short. The KDC uses an Authentication Server (i.e., Service) to handle the actual authentication process. (Baker, pp. 3-9)

The importance of Kerberos to this thesis is that it will be running natively on Microsoft's Windows 2000 products and could be a prime player in the authentication techniques used by Microsoft.

H. NAVAL AND MARINE CORPS INTRANET (NMCI)

The last concept to be covered on the background information for this thesis is the NMCI. In the last few years, the Navy, through the guidance of the Space and Naval Warfare Systems Command (SPAWAR), has been developing a plan to create the world's largest purpose-built network. This network would interconnect all Naval and Marine Corps information assets under one network. It would become a massive intranet that would span the globe, allowing over 450,000 users to exchange data from their desktops. It is projected to cost billions of dollars over a five-year period. (Peniston)

The relevance of this massive network to this thesis is two fold. First is that most of its backbone will be the public Internet itself, or at least routable to the public Internet, with the security of the data enforced through VPN connections. Naval information assets may become inter-networked over a public infrastructure,. and sites will depend on being able to create host-to-host, host-to-network, and network-to-network connections securely with VPNs. Remote access for the road-warrior will become the norm, and demand for secure remote access will only grow. This thesis will help in mapping the Navy's future implementations of remote access through VPN technology. The other relevance of this thesis to the NMCI is the hope it will aid the Naval Postgraduate School (NPS) with its transition onto the NMCI.

III. REASONS AND METHODS FOR IMPLEMENTING A VPN

A. THE NEED TO REMOTELY ACCESS THE NPS INTRANET

Certain information resources are only available from on the NPS intranet. Specifically, users gain access to programs that are loaded on machines and servers available only on the intranet, and that are protected from further distribution (i.e. copying to ones home computer) by its copyright. Files, such as a user's data, may reside on their home directories located on the central file servers, and some research resources are limited to the intranet, such as programs found at the school's library. There is also those times where a user must reside on a ".mil" domain to access certain military resources over the Internet, such as downloading the latest virus protection from the Navy INFOSEC site. Currently, it is only possible to gain access to these resources by either being directly connected to the campus intranet or dialed in to the network. This thesis recommends implementing a third option, that of a VPN for broadband users. Once connected with a VPN, all these services become available to the broadband user from his home.

The implementation of the VPN itself is really just an extension of the dial-in resources being made available to broadband users. It can also be viewed as being in direct support of the school's Policy on Appropriate Use of the Naval Postgraduate School Computing and Information Systems, NAVPGSCOL INTRUCTION 5230.4B. The purpose of this instruction is to establish the appropriate use of NPS computing and information systems, and states :

> In consideration of its primary educational mission, NPS authorizes use of its computing and information system resources for all purposes reasonably related to graduate education and research; to intellectual and scholarly inquiry; to the NPS military mission; and to the general professional interests and growth of its faculty, staff, and students. Faculty, staff, and students are encouraged to make maximum use of these resources for expanding their professional horizons, and for increasing their knowledge, skills, and ability to contribute to the NPS and to the community at large. (NAVPGSCOLINST 5230, p. 2)

So the issue becomes one of balancing the remote access needs of the faculty, staff, and students, with the security requirements imposed by those responsible for doing so in support of these resources. This becomes a sensitive issue because ease of usability and access is usually at the expense of stricter security. A primary example of this trade-off is that if our systems were not hooked up to the Internet, then they could not be exposed to hacking from outsiders, but this would be at the expense of not allowing access to all the resources that can now be found on the web. This chapter of the thesis attempts to review the options available for remote access, summarizing the advantages and disadvantages of each, in hopes that doing so will aid these decisions makers in their future decisions on remote access.

B. WHAT IS CURRENTLY AVAILABLE AT NPS FOR REMOTE INTRANET ACCESS

Before additional options can be reviewed, it is prudent to review what is currently available for remote access to the NPS intranet. This section does not necessarily deal with the broadband issues of access, but deals instead with the dial-in access being used today, and the protocols used to do so.

NPS offers a dial-up service to the faculty, staff, and students, with both a local number and a toll-free (1-800) number. These users are given a typed out procedure on how to configure their computer with a dial-up connection to the NPS Remote Access Server (RAS). The procedure describes how to configure the Microsoft Client and Dial-up Adapter, as well as mapping a drive to their home directory on the file server, and how to set up Microsoft Exchange mail client to work from their home computer.

The school has set up a modem bank of 92 modems to handle these incoming dial-up connections. These modems are 56Kbps V.90 compliant, and have a utilization rate of around 15%, with an average of 10 to 15 simultaneous connections at any given time. They authenticate users by connecting to either the Microsoft Primary Domain Controller (PDC) or one of the Backup Domain Controllers (BDC). Students are authenticated against these domain controllers with a user name and password.

The connection itself is a Microsoft version of the Point to Point Protocol (PPP). How this protocol works is the basis for many of the protocols in the next section, so it is worth further review.

1. Point to Point Protocol (PPP)

The PPP protocol is primarily used to allow communication between two computers using a serial interface, such as a computer connected to an ISP's server (Whatis?com, "Point-to-Point Protocol", 25 November 2000). It is a layered protocol, starting with a Link Control Protocol (LCP) for link establishment, configuration, and testing. This is followed by the Network Control Protocol (NCP) that is used to transport traffic between the hosts. The PPP protocol allows for the assignment of an IP address to the remote host, as well as an ability to authenticate the end user. ("Connected: An Internet Encyclopedia," 25 November 2000)

Since this section is mainly concerned with how NPS users use Microsoft's ability to connect via a PPP session, a review how Microsoft has implemented the protocol will be conducted. The following paragraphs on PPP and its four phases are taken from "Microsoft's White Paper on Virtual Private Networking: An Overview."

PPP encapsulates network protocols such as IP, IPX, and NetBEUI within PPP frames and then transmits those PPP frames across a point-to-point link. This happens in four phases.

a. Phase 1: PPP Link Establishment

PPP uses a Link Control Protocol (LCP) to establish, maintain, and end the physical connection. During this phase, basic communication options are selected, including what method of authentication will be used in phase 2. Also negotiated are what, if any, forms of compression and/or encryption will be used during the session.

b. Phase 2: User Authentication

In the second phase, a user presents their credentials to the remote access server. Most PPP implementations are limited to two types of authentication, Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP), but Microsoft has implemented and additional protocol called Microsoft Challenge Handshake Authentication Protocol (MS-CHAP).

- PAP is a simple, clear text authentication implementation. The Network Access Server (NAS) requests the user's name and password and PAP passes them in clear text form to the NAS.
- CHAP is an encrypted authentication implementation that avoids transmitting the actual clear-text password by using a hash. The NAS sends a challenge consisting of a session ID and an arbitrary challenge string to the remote host. The host must use the MD5 hashing algorithm on a concatenation of the challenge string, session ID, and the users password, and then sends that hash with the user name back to the NAS. The NAS then looks up the user's password, hashes it with the challenge and session ID and compares the hash values.
- MS-CHAP is very similar to CHAP, except it uses an MD4 hash on the challenge string, the session ID, and an already MD4 hashed password. The password is hashed on its own before being used as an input into the final hash for transport because the server does not keep a copy of the clear text password, but instead has a copy of an MD4 hash of the password. This leads to greater security, since passwords are never stored in the clear. The NAS gets its copy of the hash from either a PDC or RADIUS server. Also noteworthy is that both the client and NAS can independently generate an initial key for subsequent data encryption via Microsoft's Point-to-Point Encryption (MPPE).

c. *Phase 3: PPP Callback Control*

Microsoft's implementation of PPP includes an optional callback control phase, where after authenticating, both ends disconnect, and the NAS calls the remote host back at a specified phone number.

d. *Phase 4: Invoking Network Layer Protocol(s)*

At this point, PPP invokes various Network Control Protocols (NCPs) that were selected during the link establishment phase to configure the protocols the hosts will use to communicate. An example of these protocols might be the IP Control Protocol (IPCP) that is invoked to dynamically assign an IP address to the remote host system.

Once the four phases of the negotiation have completed, PPP begins to forward data to and from the two peers. Each packet subsequent to the negotiation is wrapped in a PPP header, which is removed by the receiving system, and decompressed and/or decrypted. The use of the PPP encapsulation then acts just like any other data link layer protocol, such as Ethernet of SONET. (Microsoft, 1999)

C. ADDITIONAL METHODS TO REMOTELY ACCESS THE INTRANET

This section covers broadband remote access methods. It is intended to serve as an introduction and overview to some of the available methods being used for remote access, not as an in-depth analysis of each. It is the author's intent that these overviews will lead to a general understanding of how the methods work, allowing the reader to pursue other resources once a choice in methodology is made.

1. Point-to-Point Tunneling Protocol (PPTP)

a. Overview

PPTP was designed to allow hosts to connect to a RAS server from anywhere on the Internet and yet still have the same authentication and access to the corporate LAN as if they were *dialing* directly into a RAS server. But instead of directly dialing the network, end users dial into their ISPs to connect to the Internet and then use PPTP to set up a secure connection to their RAS server over the Internet. PPTP then acts as a tunneling protocol, first encapsulating the network protocol datagrams (including such protocols as TCP/IP, NetBEUI, IPX/SPX) within an IP envelope, and then encapsulating that packet in a PPP envelope for transmission to the ISP. This establishes the VPN, allowing for the secure remote access via your ISP. The only difference with broadband access is that there is no need to encapsulate the IP packet within a PPP packet, since broadband utilizes an IP connection to the ISP. Also noteworthy is that PPTP was designed to use the existing PPP infrastructure, thus gaining the advantages of the PPP protocol, including dynamic address assignment from a Dynamic Host Configuration Protocol (DHCP) server residing on the secured network, user-based authentication, and compression. (Scott, Wolfe, and Erwin, pp. 62-65)

PPTP was jointly developed by Ascend Communications, U. S. Robotics, 3Com, ECI Communications, and Microsoft Corporation. Its main purpose was to provide a virtual private network between remote access users and network servers. Like other tunneling protocols, PPTP is used to tunnel PPP, an OSI layer 2 protocol, which operates at the Data Link Layer, via IP at layer 3, known as the Networking Layer. User authentication can take place via either PAP or CHAP, and encryption uses the RC4 cipher with either 40-bit or 128-bit keys. (Scott, Wolfe, and Erwin, pp. 62-65)

The PPTP encapsulation protocol is based upon the Internet standard Genetic Routing Encapsulation (GRE) protocol, detailed in RFCs 1701 and 1702. The PPTP packet is comprised of four parts: the delivery header, an IP header, a GREv2 header, and the PPP payload packet. The delivery header is the framing protocol for the medium the packet is traversing over, such as Ethernet, Frame Relay, or PPP. The IP header contains the required IP information, such as source and destination addresses, and packet length. The GREv2 header contains the information on the encapsulation method used, as well as any PPTP specific data the host or server may require. Lastly, the payload packet contains the encapsulated PPP datagram itself. (Scott, Wolfe, and Erwin, p. 70)

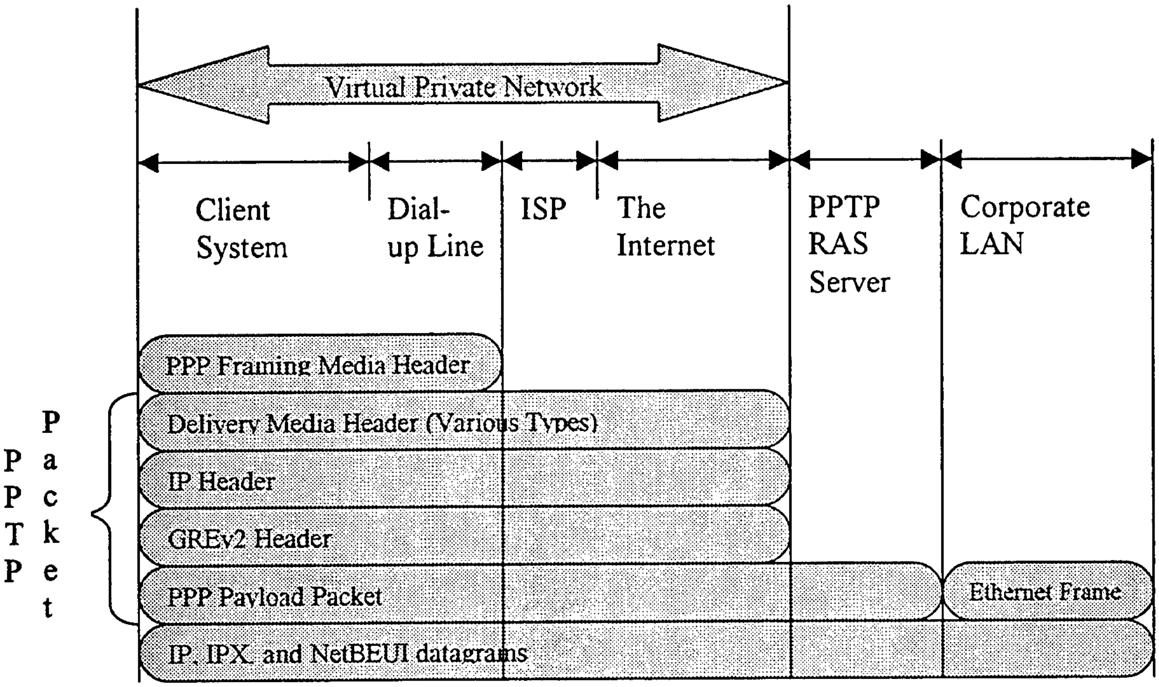

Figure 3-1: Active Protocol Layers during a PPTP Connection.
After (Scott, Wolfe, and Erwin, p. 71)

The actual encapsulation process takes place in four steps. In the first step, the end user dials in to his ISP, utilizing a PPP session. All data that is transmitted between the user and ISP will be surrounded in a PPP protocol frame. In the second step, the end user starts a PPTP connection with the RAS server that is internal to the corporate LAN, or intranet. Now the data will not only be encapsulated in a PPP protocol frame, but it will also be surrounded by the four parts of the PPTP packet describe in the previous paragraph. The third step involves the RAS server stripping off the delivery header, the IP header, the GREv2 header info, and validating the PPP client before initiating the last step of removing the PPP framing from the PPP payload packet and reformatting the packet with the appropriate medium frame for the corporate LAN. (Scott, Wolfe, and Erwin, pp. 70-71)

b. Microsoft's Implementation of PPTP

Microsoft has been the leader in deploying PPTP and because of their huge product-placement advantage, as well as the Navy's dedication to the NT platform for its IT21 standard, additional Microsoft specifics will be reviewed.

The Microsoft operating systems come with VPN capabilities inherent in the operating system. Windows 95, 98, NT and 2000 all support Microsoft's implementation of PPTP. In these versions, Microsoft has implemented the algorithms and protocols in their own unique way, and called the implementation Microsoft PPTP. The authentication protocol is the Microsoft Challenge/Reply Handshake Protocol (MS-CHAP) and the encryption protocol is called Microsoft Point-to-Point Encryption (MPPE). It should be noted that the combined protocols are referred to as MS-CHAP, and it is on its second revision, MS-CHAPv2, after receiving significant criticism of their implementation of their initial version. (Schneier and Mudge, p. 1)

As noted, Microsoft's first version of MS-CHAP received significant criticism from network security companies such as Counterpane and L0pht. These reviews left the industry with a belief that PPTP, though a solid concept, was too poorly implemented by Microsoft and should not be deployed on any network. But with MS-CHAPv2, both Counterpane and L0pht agreed that the changes did correct the major security weaknesses. (Schneier and Mudge, p. 2)

The following is the list of steps Windows uses to create a PPTP VPN, and is taken directly from Bruce Schneier and Mudge's paper on MS-CHAPv2:

> (1) Client requests a login challenge from the Server.
>
> (2) The Server sends back a 16-byte random challenge.
>
> (3a) The Client generates a random 16-byte number, called the "Peer Authenticator Challenge."
>
> (3b) The Client generates an 8-byte challenge by hashing the 16-byte challenge received in step (2), the 16-byte Peer Authenticator Challenge generated in step (3a), and the Client's username.
>
> (3c) The Client creates a 24-byte reply, using the Windows NT hash function and the 8-byte challenge generated in step (3b).
>
> (3d) The Client sends the Server the results of steps (3a) and (3c).
>
> (4a) The Server uses the hashes of the Client's password, stored in a database, to decrypt the replies. If the decrypted blocks match the challenge, the Client is authenticated.
>
> (4b) The Server uses the 16-byte Peer Authenticator Challenge from the client, as well as the Client's hashed password, to create a 20-byte "Authenticator Response."
>
> (5) The Client also computes the Authenticator Response. If the computed response matches the received response, the Server is authenticated.

The largest remaining concern for MS-CHAPv2 is the use of a client's password as an input to creating the shared secret key for the encryption that takes place between server and host. "The keys are still a function of the password, and hence contain no more entropy than the password. Even though the RC4 algorithm may theoretically have 128-bits of entropy, the actual password used for key generation have much less" (Schneier and Mudge, p. 7). Ron Cully, Microsoft's Senior Product Manager for Windows Networking, asserted that the risk of using passwords can be minimized through a sound password policy. (Raikow, p. 2)

Microsoft does not plan on incorporating any additional VPN protocols into its Windows 95, 98, or NT product lines, and will continue to support PPTP as well as introducing L2TP with IPSec in Windows 2000. Microsoft states its reasoning for continued support of PPTP is to allow the ability to create VPNs for customers who do not wish to maintain the public key infrastructure required for IPSec (Scott, Wolfe, and Erwin, p. 63).

c. *Advantages of PPTP*

A quick summary of advantages of PPTP include:

- Ease of implementation, with no additional hardware required if a company is already a Windows shop.
- No need to implement IPSec, including not having to manage certificates or shared secret keys.
- Cross platform implementations available from such third parties as Efficient Networks, Inc. These products not only allow Macintosh as

well as PC access to the network, but also improve the security found in Microsoft's PPTP.

- Since PPTP uses IP, there should not be any problems with ISP conflicts that can happen with IPSec, such as not passing IP Protocol 50 traffic.
- Changes to the existing network are minimal, including the opening of the PPTP well known port 1723, verses having to reconfigure the entire firewall with additional VPN software, or install an additional gateway appliance.
- PPTP also works exceptionally well with Network Address Translation.

d. Disadvantages of PPTP

Disadvantages to PPTP include:

- Since MS-CHAPv2 still bases encryption on a user's password, poorly implemented password policy and poor passwords can make the system more vulnerable to attack.
- Implementation does require opening an additional firewall port.
- Key length is either 40-bits or 128-bits, no longer.
- MS-CHAPv1 is still available, so a concerted effort would need to be made to ensure clients used only version 2.
- Since PPTP's primary goal is connectivity via tunneling, security is a secondary priority.

- Older Window's solutions (95/98/NT) have not scaled well to larger numbers (50+) (Brenton and Elfering, p. 37).
- Can lead to interoperability problems when used with other systems, such as may be found between different countries and different allies in a military situation.

2. Layer 2 Tunneling Protocol (L2TP)

a. Overview

L2TP is a combination of the PPTP protocol and Cisco System's Layer 2 Forwarding (L2F) protocol. It is similar to PPTP in that it relies on PPP to establish the dial-up connection, but unlike PPTP, it defines its own tunneling protocol with use of its own message types. It uses the PPP, PAP and CHAP for authentication, and because it's a layer 2 protocol, it allows for the transportation of non-IP protocols. It is also media independent, so it works over ATM, Frame Relay, or IP. (Brown, pp. 283-284)

The setup of a VPN with L2TP is also similar to how PPTP sets up. The data packet includes the PPP communications. It is also able to use PPP for encrypting the packets. (Brown, pp. 283-284)

The actual communication takes place between a L2TP Network Server (LNS) and an L2TP Access Concentrator (LAC), which allows for multiple connections inside an individual tunnel by assigning a unique Call ID to each session. L2TP then utilizes message types to communicate. Control messages are responsible for session management, including establishing and tearing down the session. These messages are also used to control the characteristics within the tunnel, including flow control, transmission rates, and buffering of the PPP packets. Data messages are the other type of messages used by L2TP and consist of the PPP packets with the framing information. (Brown, pp. 283-284)

b. Microsoft's Implementation of L2TP

As previously stated, with Microsoft being the industrial leader in home and business operating systems, and the Navy's dedication to the IT21 NT standard, Microsoft's implementation of L2TP is also worthy of review.

In its Windows 2000 product line, Microsoft has expanded its VPN implementation with the introduction of L2TP/IPSec. They have integrated the technology into Windows 2000 in order to provide an additional method of secure, low cost remote access; designing this product to inter-operate with other VPN software and devices that support Internet industry standards. In that vein, Microsoft promotes the L2TP/IPSec combination as the only standards-track technology that addresses advanced security as well as user authentication and DHCP address assignment; all other non-L2TP/IPSec implementations of user authentication and address assignments are non-standard, proprietary implementations that should be avoided. (Microsoft, *Windows 2000-Based Virtual Private Networking: Supporting VPN Interoperability,* p. 2)

IPSec tunnel mode, as defined by the standards, does not support legacy user authentication methods (PAP and CHAP), DHCP tunnel IP address assignment and configuration, and multiple protocols. So to provide truly interoperable solutions, Windows 2000 uses L2TP in combination with IPSec to provide interoperability. By placing the L2TP as a payload within an IPSec packet, communications then benefit from the standards-based encryption, integrity and protection of IPSec, while also benefiting from user authentication, tunnel address assignment and configuration, and multiple protocol support of PPP-based tunneling. (Microsoft, *Windows 2000-Based Virtual Private Networking: Supporting VPN Interoperability*, p. 3)

IPSec tunnel mode in its original specification only supports authentication via user certificates or pre-shared secret passwords. When IPSec is combined with L2TP, it allows for additional methods of user authentication. L2TP uses PPP as the method of negotiating user authentication, therefore allowing for the use of PAP, CHAP, and MS-CHAP for user authentication. It can also support advanced authentication services through Extensible Authentication Protocol (EAP), which can be used to provide plug-in authentication services within the L2TP encrypted packet within the IPSec packet. This allows for the integration of those extended authentication services such as RADIUS and LDAP via the accepted standards instead of in a proprietary way. (Microsoft, *Windows 2000-Based Virtual Private Networking: Supporting VPN Interoperability*, p. 5)

Microsoft is dedicated to only using L2TP with IPSec as the native Windows 2000 VPN IPSec solution. It will also continue to support PPTP in Windows 2000 for those implementations that require the use of non-IPSec solutions.

c. *Advantages of L2TP*

A quick summary of advantages of L2TP include:

- Ease of implementation, with no additional hardware required if a company is already a Windows shop.

- Because of its use of PPP, L2TP can authenticate users with legacy password based systems such as PAP, CHAP, and MS-CHAP, as well as advanced authentication with EAP. Previous key weakness based upon passwords is not an issue, since these packets are wrapped inside of an IPSec tunnel, using IKE for symmetric key generation.
- Changes to the existing network is minimal, including the opening of the L2TP well known port 1701, verses having to reconfigure the entire firewall with additional VPN software, or install an additional gateway appliance.
- Use of IPSec.

d. Disadvantages of L2TP

Disadvantages to L2TP include:

- Need to implement IPSec, including having to manage certificates or shared secret keys.
- Implementation does require opening an additional firewall port.
- Not compatible with network address translation (NAT).
- May not scale as well as a dedicated hardware solution such as a VPN gateway due to the intensive math required for encryption. In most VPN gateways, the hardware is custom designed to perform this function, therefore increasing speed and efficiency.

3. Secure Shell (SSH)

Though initially used by the UNIX community, SSH is making a new life for itself as it is ported to other platforms. Its original purpose was to replace unsecured protocols such as remote login (rlogin) used in TELNET, remote procedure call (rpc), and remote shell (rsh) used with FTP, which allowed a remote user to communicate and send commands back to a server. It is reviewed in this thesis because of its ability to tunnel PPP sessions to create VPNs and because of its growing acceptance as a viable option for secure remote access. (Acheson, p. 5)

a. Overview

SSH works on a PKI method where public and private keys are generated for each host/server combination. This means that in order to use SSH, each host must create a set of keys to use with each server. These keys are then used to authenticate a host each time it logs onto the server. SSH uses the RSA PKI technology to initialize a secure session, and to authenticate the user. (Acheson, pp. 6- 20)

SSH operates on well-known port TCP/22. When used as a VPN, all communication takes place on this port, and if required, is forwarded to other ports once decrypted. This is what allows the use of other ports across the tunnel, such as retrieving mail with IMAP over port TCP/143. (Brenton and Elfering, p. 69)

As stated earlier, what is making SSH a viable option for a VPN is not necessarily its inherent use on UNIX and Linux systems, but its porting to additional platforms such as Microsoft's Windows and Apple's Macintosh. A leader in this area is SSH Communications Security (www.ssh.com), and they have ported the product to

Windows via their SSH Sentinel application. The advantage of such products as these is the easy-to-use, step-by-step graphical user interface installation wizards that are used to create keys and configure the security policy. Also noteworthy is the ability to operate with other vendor's IPSec gateways and implementations of VPN products.

b. Advantages of SSH

A quick summary of advantages of SSH include:

- User friendly set-up interface.
- Only well known port 22 and IP is required to be supported by the firewall.
- Cross platform implementations, including www.ssh.com for Windows and www.macssh.com for Macintosh.
- Depending upon the implementation, SSH products can be very inexpensive (even free). (Brenton and Elfering, p. 70)

c. Disadvantages of SSH

Disadvantages to SSH include:

- Since authentication is based upon a PKI implementation, key management can become cumbersome when keys must be generated for every server/host combination. (Acheson, p. 55)
- The preferred version of SSH is version 2, which is still being developed by an IETF work group. Without this finalized standard, some implementations are considered buggy. (Acheson, p. 54)

- Though PPP over SSH is a viable VPN, other implementations such as VPN gateways are usually a more robust VPN solution.

4. TELNET

Though considered by some a viable option for remote access, the author of this thesis concurs with Steven Brown's observation that TELNET was designed for general, bi-directional, 8-bit, byte-orientated *unsecured* communications. Security considerations were never built into TELNET. It should continue to be used in the context for which it was designed: simple communication. As such, it will not be covered in any more depth in this thesis. (Brown, pp. 461-462)

If TELNET is a reader's desired remote access method, SSH should be considered its proper replacement.

5. VPN Gateway Appliances

The last category of access methods is the VPN gateway appliance. This is a dedicated hardware and/or software solution that resides on the boundary of network and is responsible for VPN activities. It can be either integrated via software into a computer that already resides on the network, such as in a firewall, or can be a completely separate appliance that runs its own operating system. An example of the later solution would be a VPN TimeStep Gateway from Alcatel.

There are advantages and disadvantages to either of these solutions and industry experts argue over the merits of each. Some prefer to place more of a burden on a component that already exists on the network, such as enabling the VPN software found in some firewalls, then to install an additional component on the network. This is so that you do not have to add another piece of equipment that must be managed and because it does not introduce one more potential access point for a hacker. But others prefer to layer security with multiple components to keep from having a single point of failure, and that it is better not to over burden a single piece of equipment. There are also those who question how strong a product is if it tries to do too much and be an all in one solution. Throughout the author's research, there seemed to be no definitive answers to these issues, only many opinions. Two points did seem to be generally agreed upon though. The first was to ensure that if you do use an all in one solution, be sure that it does not overtax the processing power of that box (i.e. greater then 85-90% of CPU utilization), especially since the cryptography function can be extremely math intensive. The second was that using a hardware solution for the cryptography, such as found in a VPN gateway by Intel or Alcatel, can increase the speed and efficiency of the sessions, since the encryption is done by dedicated hardware, not via the software in the OS.

Whichever gateway appliance implementation is chosen doesn't matter for this review though, since both usually have the same options for how they are configured, including the choice of whether to implement IPSec with either a shared secret password authentication or through a PKI with a certificate authentication.

a. IPSec with Shared Secret Authentication

VPN appliances can be configured with a shared secret password to authenticate a host or person. Remember that VPNs can be established between two gateways, so it is not always a person who must be authenticated, and that IPSec was originally designed to authenticate the host. Therefore, some implementations compliment the initial authentication via a shared secret password with an extended authentication mechanism such as RADIUS. This use of XAUTH can either give an additional layer of authentication, or can be the primary authentication when a common shared secret password is used. What is meant here is that some organizations configure their VPN client-ware with an embedded shared secret password, making the software the first part of a "what you have, what you know" security combination. The second part is then the extended authentication of a user name and password via an authentication server, such as a RADIUS server. This enables the VPN appliance to be configured with only one password making it much easier to manage, and yet still ensures security, and authentication via RADIUS. Though this method may seem a bit unorthodox, it can be a very viable method for an implementation, as has been proven by Bentley College.

It was during the 1999 PKI and VPN Workshop that Captain Galik, Program Manager for Navy Information Systems Security, PMW-161 at SPAWAR, touted Bentley's remote access implementation. So on his recommendation, the author visited Bentley College to discuss their implementation and architecture. Bentley's IT department had chosen to configure their Intel/Shiva VPN Gateway in this manner, and have found it very effective for over 750 initial test cases, with no security breaches since its introduction over a year and half ago. (Interview with Cekanavich)

The management of the shared secret password becomes the main issue with this type of implementation. If an organization chooses to use individual passwords for each user, the VPN appliance becomes an additional piece of equipment that must be managed by an administrator, and users are generally required to remember another password. This method, though more management intensive, allows for the logging of individuals. It also makes it easier if a password has been compromised; only one new password must be distributed to a single user. If a common shared secret password is used, *every* piece of client-ware must be updated with the new password, and the administrator must find a way to get that new password out to all the users in a secure manner, so it is not intercepted. This can be a major hassle, easily surpassing the initial time saved by using only one password.

The author has been involved with many discussions on this topic of a common shared secret key compromise, and it leads to an interesting side note to the idea. If you are only using the shared secret password to gain access to the user authentication part of the process, is the password any different then the phone number a person would use to access a PPP connection. Authentication takes place after the initial tunnel is set up with IKE, but before the final SA is established. The password has no effect on the final key chosen, since it is created in IKE Quick Mode. So if the password is compromised, or even freely distributed, is it *really* any different than publicizing the phone number for a PPP connection? If you trust your Dial-in PPP connection to authenticate your users with extended authentication mechanisms such as an NT domain controller, or RADIUS as most ISP's do, then shouldn't that be enough for the same level of access with a different medium, that being broadband? The answer is really up to the security manager of the organization, but it is important that these security managers consider the possible merits of the argument.

b. IPSec with Certificate Authentication

This is the industry's preferred method of implementation. It is management intensive, but also the most robust in its ability to authenticate a user. Unlike passwords that may be poorly chosen and easily guessed or gotten through social engineering, compromising a certificate requires quite a bit more effort. To compromise a certificate takes reverse engineering of the owner's private key, which is on the order of 1024 bits, vs. 64 bits found in a 8 byte password. An imposing task to be sure. But this security comes at the cost of extensive management of the certificates. This can be done in house with an organization's own certificate server, or contracted out to a third party, such as VeriSign.

Certificates are not without their own problems as well. There are the issues of maintaining Certificate Revocation Lists (CRLs), used to publicize an invalid certificate, and choosing how CRLs should be used. Are they updated daily and pushed to certain servers that will be used for a validity check? Do you sacrifice speed of establishing the tunnel so each certificate can be checked against a live CRL at the CA? Key escrow is also an issue. If a person chooses to use their keys for encrypting data, and that person later leaves the organization due to unforeseen circumstances such as death, then how does an organization decrypt what might be vital company data? Does the company choose to escrow the private key, and if so, can that key really be considered private since there are multiple copies? There is also the issue of how a user stores the private key. It could be kept on a specific computer, leading to the question of how does one use multiple computers then, or is it kept on a disk which may be lost or stolen? Cost can also be an issue, since using a third party can be very expensive. The list goes on, with new issues and answers developing all the time, so it must be left up to the individual organization on how best to handle their own situation.

Either way, with shared secret passwords or certificates, the industry is choosing to standardize on IPSec. Like any other maturing standards, there is a period of growth where issues are discovered and solutions developed on how to best meet them. For now, organizations that choose to adopt technology in its earlier stages, such as VPNs and IPSec, must be willing to work with emerging standards.

c. Placement of the VPN Appliance Relative to a Firewall

When working with a VPN appliance, decisions must be made where to place it on the border of the network, specifically with regards to the firewall. Issues surround each of the following four options, and these will need to be considered by the organization before a final decision can be made as to the type and location of the VPN it chooses to deploy.

The first option is to place the VPN within the firewall itself. Many of today's firewall products come with VPN modules that can be installed into the existing firewall. As touched upon earlier, this has the advantages of not requiring addition equipment and management of an additional box, but may have the disadvantages of overtaxing the current system the firewall resides on. This option also depends upon the security manager's view on whether it is better to keep the network access limited to as few physical point of entry as possible, or it is better to layer the network with an additional security box by installing a separate appliance. This last argument seems to be a glass is half-empty or half-full issue, being based upon a manager's preference more then any evidence that one is better than the other.

The second option is to place a separate VPN appliance in front of the firewall, between the firewall and the Internet itself. This allows for any traffic that is decrypted through the VPN to be checked against the firewall rule base, but also leads to certain problems. If you are issuing network IP addresses to your remote VPN clients, which most remote VPN clients require, the firewall rule base must be changed to break one of the first rules found in a firewall, that internal addresses coming from outside the network should not be allowed to pass. This option also means your VPN may also have to act as a router, passing all non-VPN data to the firewall.

The next option is to place the VPN behind the firewall. This limits the firewall's ability to check the packets against its rule base, since packets destined for the VPN are encrypted. This can also be an issue if the firewall does not support IP Protocol 50, ESP, which is its own protocol, not TCP or UDP. If this is the case, the firewall will drop the packets instead of being able to pass them. Additional ports will have to open on the firewall as well. IKE requires UDP port 500, and if certificates are used on an LDAP server, TCP port 389 will also need to be accessible.

Seeing as how in the last option the firewall has a very limited ability to check the VPN packets, some organizations have chosen to place the VPN appliance in parallel to the firewall, allowing the screening router to decide which traffic goes to which appliance. This overcomes the issues with the placing the VPN either before or after the firewall, but still leaves the problem that the traffic is never checked against the firewall rule base. In answer to that concern, some security managers believe that since the VPN authenticates the users and treats them as if they were already on the network, a firewall is not required to protect the network from these outsiders, since they really are insiders anyway. Other security managers are placing an additional firewall behind the VPN, usually with a smaller rule base since access is limited to users who are already authenticated. It should be noted though that this additional firewall means additional cost and management, so it may not be right for an organization with limited resources.

D. ADDITIONAL ISSUES SURROUNDING VPN IMPLEMENTATION

This section is included to introduce issues that surround a VPN implementation that have not yet been covered in the previous sections. It will discuss many of the issues that have been discovered by those organizations that have been the early adopters, and who have "bled" by being on the cutting edge. The real value of this section is that it introduces those issues that have not been well publicized by the industry.

1. Problems with Internet Service Providers (ISP)

VPNs are a great tool for accessing that internal network from home, but many of today's ISPs don't work well with VPNs in general and IPSec specifically. Some ISPs who are offering broadband access in the form of cable modems or DSL have found that with the advent of VPNs, more of their home users are utilizing their Internet access for telecommuting. These selected ISPs are no longer allowing home users to pay the residential fee for use of the ISP services if they utilize VPNs. Instead, they are making these home users subscribe to a small office home office (SOHO) rate, which is usually quite a bit more money. Other ISPs just will not allow *any* VPN IPSec traffic. The point here is that before a company chooses to establish VPN remote access, it should research what ISPs its employees are using, and inform them if they may need to choose another provider.

2. Network Address Translation (NAT) for the Home User

NAT is a well-documented problem when used with AH, since the changing of the IP address invalidates the integrity check. However, it is not a problem with ESP, which is the most frequent IPSec protocol choice. Other issues with NAT are starting to arise though. Many home users are now trying to find ways to share their broadband home access among multiple machines in their home, so the spouse and kids can surf on their machines while the other host is used for remote access. The problem lies in how the home users implement NAT. In addition, some software implementations and home routers/firewalls do not work well with IPSec or VPNs. Since these products are constantly changing, specifics can not be given here, but the issue should be researched by an organization so a home user will know what products are interoperable with NAT and the organization's VPN.

3. Ensuring that Virus Protection and Home Firewalls are Used

Once a home computer connects to an intranet via a VPN, it becomes part of that network. If that home computer is infected with a virus, or has a Trojan horse residing on it, it may make the rest of the network vulnerable to attack. For instance, some Trojan horses allow that attacker to take control of a remote host that has been infected. If that host connects to the intranet, and is able to be controlled by an attacker, then that attacker has access to the intranet as well. This is a very real problem, but one that can be reduced to an acceptable level of risk. What most companies, including the Navy, are doing is obtaining site licenses for virus protection programs that include extending the license to the home systems of their employees. The company then encourages, or even requires, its home users to use this software in order to be able to access the network. This idea could, and should, be extended to home firewall products as well. Companies just need to ensure that whatever products are chosen, they are compatible with the VPN (see previous section on NAT for more information on this concept).

4. Split Tunneling

Some VPN products allow the client-ware to set up a split tunnel. What this does is force only the information that is going to the intranet to be encrypted and sent to the intranet, while normal Internet traffic is sent directly to the ISP to be routed as normal. The concept is desirable, but can be complicated to configure.

Split tunnels allow users to make a VPN connection to the intranet and perform intranet activity, while Internet activity is directed through the ISP, all while still logged onto the VPN. The alternative, not utilizing split tunneling, also allows the user to access the intranet while logged into the VPN client, but if they want to access the Internet, they must either route the Internet traffic *through* the intranet, or they must log off the VPN and directly connect to the Internet via their ISP. The concern from most network managers with the alternative is that the home user will log on to the VPN to work, and then not log off when done. This then forces all subsequent personal Internet traffic from that host through the internal network, using precious bandwidth when there is no need to do so.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. CONCLUSIONS

A. GENERAL COMMENTS

Every command should be preparing for the inevitable demand for broadband access to their intranet. Cable modems and DSL, as well as high-speed wireless access, will only proliferate in the coming years. The growing need for real time data, especially while away from the "office," will continue to place a greater demand on network manager's to allow access to their intranets.

Security of the connection and the data will be at the forefront of the network managers concerns. A paradigm shift should be taking place now, understanding that trying to secure the "plumbing" of the network becomes unnecessary when data integrity and confidentiality are built into the network at the data layer. By securing IP traffic with a VPN, there is no longer a need to pay exorbitant prices for dedicated lines, or for other means of secure plumbing. Standardizing on IP, with security at the data layer with VPN technology, will allow for faster adoption of a secure network technology that can be used for everything from data traffic to voice over IP. As IP has become the standard for the Internet, VPN technologies should become the standard for securing data that traverses a network.

Because VPN protocols are still in the early phases of design and adoption, managers will have to choose which protocols are best for their organization. IPSec is prevailing as the optimum solution, so organizations should plan on implementing it, even if other protocols are to be used earlier in the implementation. This means that PKI and certificate management should be planned for. The Navy is currently planning for and deploying a proprietary PKI solution, called DoD PKI. It is projected that all end users of the Navy will have a certificate by October 2001 (Galik, p. 15).

B. RECOMMENDATIONS FOR THE NAVAL POSTGRADUATE SCHOOL

The author's recommendation for the NPS is that the time is right to implement a VPN for broadband access to the intranet. The school can implement the VPN in a number of ways, since it is a Microsoft shop, has a firewall capable of being upgraded with its own VPN solution, and it has been given a VPN appliance from SPAWAR. The school will have to evaluate the different options and decide which is the best fit with their network topology, performance requirements, and security perspective. Which ever is chosen, a pool of addresses will have to be set aside for DHCP assignment to the remote clients, which shouldn't be a problem for the NPS since it owns the class B license of 131.120.X.X. Also, once a chosen method is implemented, NPS needs to ensure a standardized configuration for clients is published, as has been done for the remote dial-in users with the "NPS Dial-Up E-Mail Remote Access Services (RAS)" instruction sheet.

The assumptions upon which the author based these recommendations involve the overall security posture of the NPS campus. The school resides upon an open campus, with no restriction for access and no guards on the gates. Most labs remain open during working hours, so there is no restricted access to some of the individual computers on the campus for most of the day. Intranet access control is therefore based upon the authentication mechanisms in Microsoft's domain controllers and the built-in security of the operating system. Likewise, the domain controllers also control the PPP dial-in access for those currently connecting to the intranet remotely. Though it will be up to the network and security managers on the campus to decide if any of the following VPN implementations meet the current requirements for security, the author believes his recommendations are at least as stringent as what is currently in place, if not more secure.

1. A Microsoft Solution

A viable option for the NPS is to implement Microsoft's VPN technology. A Windows 2000 Server could be installed and configured to support PPTP connections from older Windows 95, 98, and NT clients and L2TP/IPSec connections from the latest 2000 product line. Actions should be taken to ensure clients only use version 2 of the PPTP for the enhanced security reasons mentioned earlier. Authentication will take place against the Window's domain controllers, just as is currently done for the PPP dial-in connections.

NPS' strong password policy and their continuous checking for weak passwords minimize the industry's concern that PPTP's use of passwords is unsuitable for key generation due to poorly chosen passwords. Because NPS requires passwords of at least 8 characters in length, with upper and lower case letters combined with digits, the passwords used by NPS members are fairly hard to guess, therefore the keys generated from them are fairly secure.

a. *Advantages of a Microsoft Solution*

A quick summary of the advantages of a Microsoft solution include:

- A quick and simple setup that would not require any additional hardware aside from installing Windows 2000 Server on an available computer.
- Use of PPTP does not require the implementation of IPSec. Authentication is done with the use of the existing domain controllers.
- IPSec can be incorporated with Window's L2TP/IPSec on a trial basis with Window's 2000 clients when the school is ready to implement IPSec.
- No additional client-ware is required for personnel who use the Windows operating system.

b. *Disadvantages of a Microsoft Solution*

Disadvantages of a Microsoft solution include:

- Does not support other operating systems, such as Apple Macintosh and Linux, both of which are prevalent in an academic environment such as at the NPS.
- As stated earlier, may not scale well to more then 50 simultaneous connections.
- May be slower than a hardware appliance solution where the encryption takes place at the hardware level.

- May not be interoperable with other VPNs, such as can be found on the NMCI.

An excellent white paper can be found on the Microsoft web site that contains step-by-step instructions on the implementation of the Microsoft solution. It's entitled "Windows 2000 Virtual Private Networking Scenario."

2. Upgrading the Current Firewall Product with a VPN Module

During Fiscal Year (FY) 2000, the NPS integrated a new firewall product onto the campus intranet. This new firewall was produced by a company call Raptor Inc. It has the ability to be upgraded with a VPN solution that uses a client called Mobile. Though not a primary consideration in the purchasing decision, it was hoped that the Mobile client could be implemented in the future.

During the last year, a company called Axent bought out Raptor Inc. With this purchase, Axent introduced a new VPN solution called PowerVPN, which is a stand-alone, firewall-independent VPN. With this introduction, Axent has limited its support for the Raptor upgraded firewall VPN product and for the Mobile client-ware. (Axent Inc.)

This development makes the use of Mobile client a less attractive alternative, though still a viable one.

a. Advantages of a Upgrading the Current Firewall

A summary of the advantages of upgrading the current firewall include:

- As an integrated component, it is designed to work well with the current network firewall.

- No need to purchase additional equipment.
- No requirement to manage an additional network device.
- Supports password, NT domain controller, and RADIUS authentication (Raptor Inc.).

b. Disadvantages of a Upgrading the Current Firewall

Disadvantages of upgrading the current firewall include:

- Known not to scale well beyond 5 or 10 simultaneous connections (Axent Inc.).
- Limited support for the product from the parent company.
- Limited support for different platforms and equipment (Raptor Inc.).
- Requires the cost of obtaining and installing the upgrade.
- Limited support for stronger authentication methods (Raptor Inc.).

3. Installation of a VPN Appliance

The third option is to install a separate VPN appliance, such as Alcatel's TimeStep VPN appliance, or Axent's PowerVPN software. This option is by far the most complex to install and configure, but it also gives the best performance, scalability, and flexibility. It is also this author's recommended solution.

This option is favorable over the others for its flexibility and because it will ensure interoperability with other components of the NMCI by using the Navy's chosen product, the Alcatel TimeStep solution. This product not only ensures interoperability, it also ensure compliance with the DoD PKI model, since Alcatel has committed to supporting that proprietary implementation of the PKI. In this specific instance, it also saves NPS the cost of the appliance, since SPAWAR was generous enough to donate a TimeStep 7520 box to the school, at a cost of over $10,000.

The 7520 model supports 70 Mbps of bandwidth, and up to 2000 simultaneous connections. The client software runs on all the Window's operating systems, and originally ran on the Mac OS as well, though there is questions as to whether or not they will continue support for the Mac. Even if they don't, third party software developers such as Netlock (www.netlock.com) have created interoperable client-ware that will work with TimeStep. Alcatel's solution even supports SWAN, an IPSec and IKE implementation for Linux. A site license can also be obtained to freely distribute the TimeStep client-ware.

Until individual certificates are issued by the Navy and DoD PKI is finalized and supported by Alcatel, the author recommends using the XAUTH user authentication method. As discussed earlier, NPS could set up a single RADIUS server that would authenticate clients with the Microsoft domain controllers. Then the TimeStep box could be configured to use XAUTH with the RADIUS server. The client-ware could be configured with an encrypted IPSec shared secret password, and freely distributed to NPS personnel, as was discussed earlier with Bentley College's implementation. This configuration equates to the PPP dial-in method of remote access, with the shared secret password gaining the same access to the intranet as the dial-in phone number, with both methods then using the domain controllers for authentication of the individual. This implementation is better than the PPP dial-in method since it ensures security through data encryption, tunnel security through IKE key generation, and supports broadband access.

If this method of freely distributing a common IPSec shared secret password is not a viable option to the network security manager, then the author recommends the use of individually assigned passwords. The TimeStep box has the ability to assign passwords to individuals for the IPSec authentication section, and then continue to use XAUTH for user authentication. This method requires more management for the administrator of the TimeStep box. If this method is chosen, it may be better not to distribute the client-ware to all personnel, but to issue it on a case by case basis.

In either case, NPS will still have the ability to transition over to certificates once DoD PKI is in place. The TimeStep box can even be used with both types of IPSec authentication at the same time, to ease the transition by migrating just a few users at a time.

The author also recommends contacting SPAWAR San Diego for assistance in designing and implementing their VPN. The team there has dedicated themselves to designing extranets with VPN technology, and are experts in the field. Another team of experts can be found at the Marine Corps Network Operations Center (NOC) in Quantico, Virginia. This team has already installed these TimeStep VPNs at over 27 Marine Corps bases, creating one large extranet between them. Alcatel also has a support team dedicated to the DoD, and can be contacted via email at usdod.support@alcatel.com.

a. *Advantages of a Using a VPN Appliance*

A summary of the advantages of using a VPN appliance include:

- A dedicated appliance allows for the greatest level of network design flexibility.
- A hardware appliance, such as a TimeStep gateway, scales well up to 2000 simultaneous connections.
- It allows for an IPSec transition from a shared secret password authentication to the use of DoD PKI when it becomes available.
- For the NPS, SPAWAR's donation of the TimeStep 7520 allows for significant cost savings over installing another type of appliance or firewall upgrade, and ensures NMCI interoperability.
- With encryption and decryption at the hardware level, as is done by the TimeStep gateway, there will be a significant increase in the speed of the connection without overtaxing the gateway, as may happen with an upgraded firewall or Microsoft only solution.
- With the use of a RADIUS server authenticating to the domain controllers, the complexity of implementing an authentication mechanism is kept to a minimum.
- Leverages the USN/USMC experience with deployment and maintenance issues.
- OS independent, since it has the ability to support Windows, Macintosh, and the UNIX/Linux platforms.

b. Disadvantages of a Using a VPN Appliance

A summary of the disadvantages of using a VPN appliance include:

- It is the most complex implementation of the available choices.
- There is increased management with the use of a shared secret IPSec authentication method if the shared secret is not freely distributed, or if individual passwords are managed on the gateway.

C. THESIS APPLICATION TO OTHER MILITARY ORGANIZATIONS

This thesis culminates over a year's worth of research, including interactions with many of the known experts in the field. It should assist readers in building the necessary background for designing their own implementation by introducing them to the basic concepts and the pitfalls that have been discovered over the last few years.

With security policy being open to interpretation, it is not possible to make specific network design recommendations. This thesis though will allow other commands to understand how VPN technology works, and how to design an implementation based upon their own requirements.

It is the author's hope that this thesis will be of value to those military organizations that have an interest in creating a remote access mechanism to their intranets. It could be used by any command as a starting point to designing their implementation, introducing the concepts and bringing forward those issues of security and protocol choices that will need to be selected for their specific implementation.

There are a great number of additional resources that can be consulted when designing a VPN implementation. Some have already been mentioned in the last paragraph of the previous section, which included SPAWAR, USMC NOC, and Alcatel. Other resources of course include those found under the References section of the thesis. Also available are such organizations as SANS (www.sans.org) and VPN Conference and Exhibits (www.vpncon.com) who offer excellent conferences and specific tracts on how to design and implement a VPN. The VPN Consortium (www.vpnc.org) is also an excellent place to learn about the latest issue concerning VPNs, including RFCs and other current issues.

D. FUTURE AREAS OF STUDY

While researching this thesis, the author discovered other collateral areas that would be excellent topics for future thesis research. These include:

- A thorough look into the Navy's security policies in light of VPN technologies. VPNs encompass many new protocols and network design issues that need to be evaluated in conjunction with security policy. A Navy-wide policy recommendation could be the culmination of the research.
- VPN vulnerabilities need to be researched in conjunction with remote access, specifically how to best prevent subversion of remote systems, such as home users.
- The implications IPv6 will have on VPN implementation.

E. FINAL COMMENTS

This final section is dedicated to some final comments that the author hopes will help shift some paradigms. The first paradigm is one of bits and security. At a recent conference, the author attended a panel discussion where a gentleman stated "Why is it that I can trust my entire life's savings to a 4 digit PIN number with my bank card, but I need 1024 bits to check my email securely?" Though there are a lot of reasons based upon policy (e.g., limited to $500/day withdrawal) and a bunch of other reasons discussed by the panel, it does make one question how much security really is needed? As far as the author is concerned, too many people are too concerned with bit length of keys and implementations of PPTP, and not enough about what is the *value of the data* that is trying to be protected. In an open campus environment, where a hacker could walk right into a school lab, and modems proliferate throughout the campus, is IPSec *really* required?

The second paradigm deals with "securing the plumbing." So much time, energy, and money is spent trying to ensure that no one can gain access to the network, tap the cable, that we lose the idea that if encryption is at the data layer, securing the plumbing becomes unnecessary. This concept will become even more important to understand when we consider that the age of wireless communications, from phones to broadband, will make securing the medium impossible. It needs to be understood that the future lies in encrypting objects, such as the datagrams that traverse a VPN, not in securing the plumbing. If this concept is combined with a standardization on IP as the protocol, the paradigm can truly shift for the DoD to an understanding that the future can lie in all data traversing IP with security at the data level, for everything from voice over IP to computer data. The value of the data then dictates the security required, and the key length that will be needed.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

Acheson, S., "SSH: Introduction Through Implementation," SANS Sixth Annual Conference on Securing Networks and Systems, 15 October 2000.

Axent Inc., "Axent to Release New VPN, Firewall Products," [http://www.internetwk.com/story/INW20000107S0006], 6 December 2000.

Baker, M., "UNIX@NIGHT: Kerberos," SANS Sixth Annual Conference on Securing Networks and Systems, 20 October 2000.

Bourne, T., Gaidosch, T., Kunzinger, C. Murhammer, M., Weinfurter, A., "A Guide to Virtual Private Networks," Prentice Hall PTR, 1998.

Brenton, C., and Elfering, D., "VPNs and Remote Access," SANS Parliament Hill 2000, 24 August, 2000.

Brown, S., "Implementing Virtual Private Networks," McGraw-Hill, 1999.

Erwin, M., Scott, C., and Wolfe, P., "Virtual Private Networks,"2nd Edition, O'Reilly, 1999.

Galik, D., "PKI and the Navy," *CHIPS*, Volume XVIII Issue I, Winter 2000.

Interview between Alan Cekanavich, Director, Systems, Networks and Telecommunnications,*Bentely College*; Ed Borden, Senior Network Analyst, Systems and Networks, *Bentley College*; and the author, 15 September, 2000.

Maier, P., "Designing and Building Extranets: Step-by-Step," SANS Sixth Annual Conference on Securing Networks and Systems, 21 October 2000.

Microsoft, White Paper, *Virtual Private Networking: An Overview,* 1999.

Microsoft, White Paper, *Windows 2000-Based Virtual Private Networking: Supporting VPN Interoperability,* 2000.

Naval Postgraduate School, NAVPGSCOLINST 5230, "Policy on Appropriate use of Naval Postgraduate School Computing and Information Systems," 19 February 1999.

Newbridge, "Enabling Managed Secure VPN Services." [http://www.timestep.com/doctypes/solutionsheet/timestep/growth.jhtml]

Peniston, B., "Navy-wide intranet to be online by 2001," The Navy Times, 19 July 1999.

Raikow, D., "Windows-Based VPNs Not 'Industrial Strength'?" [http://www.zdnete.com/filters/printerfriendly/0,6061,2293711-79,00.html]. 26 November, 2000.

Raptor Inc., "EagleMobile 4.1 Frequently Asked Questions (FAQ)," [http://www.raptor.com/products/faqs/mobile.html], 6 December 2000.

Schneier, B., and Mudge, "Cryotanalysis of Microsoft's PPTP Authentication Extensions (MS-CHAPv2)." [http://www.counterpane.com/pptpv2-paper.html]. 26 November, 2000.

Snyder, J., "Introduction to VPNs and Advanced VPN Training," *Tutorial Proceedings, Virtual Private Networks Conference & Exhibition,* Boston, MA, 11 September, 2000.

TimeStep, *Understanding the IPSec protocol suite*, 1998.

Walton, C., "IPv6 At the Starting Line," NetWare Connection, May Issue, Volume 10, Number 5, 1999.

WhatIs?com, "Lightweight Directory Access Protocol." [http://www.whatis.com/WhatIs_Definition_Page/0,4152,214076,00.html]. November, 2000.

WhatIs?com, "Point-to-Point Protocol." [http://whatis.techtarget.com/WhatIs_Definition_Page/0,4152,214311,00.html]. 25 November, 2000.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center.. 2
8725 John J. Kingman Road, STE 0944
Fort Belvoir, VA 22060-6218

2. Dudley Knox Library.. 2
Naval Postgraduate School
411 Dyer Road
Monterey, CA 93943-5101

3. LT Richard Scott Cote'.. 1
1165 Third St.
Monterey, CA 93940

4. Professor Daniel Warren, Code CS/Wd 1
Naval Postgraduate School
Monterey, CA 93943-5118

5. Professor Rex Buddenberg, Code IS/Bu.. 1
Naval Postgraduate School
Monterey, CA 93943-5118

6. CO, SPAWARSYSCEN San Diego.. 1
53560 Hull St.
San Diego, CA 92152-5001
Attn: John F. Feist
Bldg. 40, rm 30

7. Director, MITNOC.. 1
Attn: GySgt Michael R. Davis
Logistics Branch, 3086 Roan Street
Quantico, VA 22134

8. Ian Hamilton and Stephane Durette 1
Tower 1, 5th floor
600 March Rd.
Kanata, Ontario
Canada, K2K 2E6

9. Chair, IS Academic Group, Code IS 1
Naval Postgraduate School
Monterey, CA 93943-5118